

# The Projected Power Method: An Efficient Algorithm for Joint Alignment from Pairwise Differences

Yuxin Chen <sup>\*</sup>      Emmanuel J. Candès <sup>\*†</sup>

September 20, 2016

## Abstract

Various applications involve assigning discrete label values to a collection of objects based on some noisy data. Due to the discrete—and hence nonconvex—structure of the problem, computing the maximum likelihood estimates (MLE) becomes intractable at first sight. This paper makes progress towards efficient computation of the MLE by focusing on a concrete joint alignment problem—that is, the problem of recovering  $n$  discrete variables  $x_i \in \{1, \dots, m\}$ ,  $1 \leq i \leq n$  given noisy observations of their modulo differences  $\{x_i - x_j \bmod m\}$ . We propose a novel low-complexity procedure, which operates in a lifted space by representing distinct label values in orthogonal directions, and which attempts to optimize quadratic functions over hyper cubes. Starting with a first guess computed via a special method, the algorithm successively refines the iterates via projected power iterations. We prove that the proposed projected power method makes no error—and hence converges to the MLE—in a suitable regime. Numerical experiments have been carried out on both synthetic and real data to demonstrate the practicality of our algorithm. We expect this algorithmic framework to be effective for a broad range of discrete assignment problems.

## 1 Introduction

### 1.1 Nonconvex optimization

Nonconvex optimization permeates almost all fields of science and engineering applications. For instance, consider the structured recovery problem where one would like to recover some structured inputs  $\mathbf{x} = [x_i]_{1 \leq i \leq n}$  from noisy samples  $\mathbf{y}$ . Finding the maximum likelihood estimates (MLE)

$$\text{maximize}_{\mathbf{z} \in \mathbb{R}^n} \quad \ell(\mathbf{z}; \mathbf{y}) \quad \text{subject to} \quad \mathbf{z} \in \mathcal{S} \quad (1)$$

could often be highly nonconvex, depending on the choices of the log-likelihood function  $\ell(\cdot)$  as well as the feasible set  $\mathcal{S}$ . In contrast to convex optimization, which has become the cornerstone of modern algorithm design, nonconvex problems are in general daunting to solve. Part of the challenges arises from the existence of (possibly exponentially many) local stationary points; in fact, oftentimes even checking local optimality for a feasible point proves NP-hard.

Despite the general intractability, recent years have seen progress on nonconvex procedures for several classes of problems, including low-rank matrix recovery [1–10], phase retrieval [11–18], dictionary learning [19, 20], blind deconvolution [21], empirical risk minimization [22], to name just a few. For example, we have learned that several problems of this kind provably enjoy benign geometric structure when the sample complexity is sufficiently large, in the sense that all local stationary points (except for the global optimum) become saddle points and are not difficult to escape [6, 10, 19, 20, 23]. For the problem of solving certain random systems of quadratic equations, this phenomenon arises as long as the number of equations or sample size exceed the order of  $n \log^3 n$ , with  $n$  denoting the number of unknowns [16].<sup>1</sup> We have also

<sup>\*</sup>Department of Statistics, Stanford University, Stanford, CA 94305, U.S.A.

<sup>†</sup>Department of Mathematics, Stanford University, Stanford, CA 94305, U.S.A.

<sup>1</sup>This geometric property alone is not sufficient to ensure rapid convergence of an algorithm.

learned that it is possible to minimize certain non-convex random functionals—closely associated with the famous phase retrieval problem—even when there may be multiple local minima [12, 14]. In such problems, one can find a reasonably large basin of attraction around the global solution, in which a first-order method converges geometrically fast. More importantly, the existence of such a basin is often guaranteed even in the most challenging regime with minimal sample complexity, which, for example, is the case in phase retrieval when the sample size is on the order of  $n$  [14]. This motivates the development of an efficient two-stage paradigm that consists of a carefully-designed initialization scheme to enter the basin, followed by an iterative refinement procedure that is expected to converge within a logarithmic number of iterations [12, 14]; see also [1, 2] for related ideas.

In the present work, we extend the knowledge of nonconvex optimization by studying a class of assignment problems in which each  $x_i$  is represented on a finite alphabet, as detailed in the next subsection. Unlike the aforementioned problems like phase retrieval which are inherently continuous in nature, in this work we are preoccupied with an input space that is discrete and already nonconvex to start with. We would like to contribute to understanding what is possible to solve in this setting.

## 1.2 A joint alignment problem

This paper primarily focuses on the following joint discrete alignment problem. Consider a collection of  $n$  variables  $\{x_i\}_{1 \leq i \leq n}$ , where each variable can take  $m$  different possible values, namely,  $x_i \in [m] := \{1, \dots, m\}$ . Imagine we obtain independent pairwise difference samples  $\{y_{i,j} \mid (i, j) \in \Omega\}$  over some symmetric<sup>2</sup> index set  $\Omega \subseteq [n] \times [n]$ , where  $y_{i,j}$  measures the modulo difference of the incident variables  $x_i$  and  $x_j$ ; more precisely,

$$y_{i,j} \stackrel{\text{ind.}}{=} x_i - x_j + \eta_{i,j} \bmod m, \quad (i, j) \in \Omega, \quad (2)$$

where the additive noise  $\{\eta_{i,j} \mid i > j\}$  are i.i.d. random variables supported on  $[m]$ , and we take  $\eta_{j,i} = -\eta_{i,j}$  for all  $(i, j) \in \Omega$ . Here and below,  $\Omega$  is obtained via random sampling at an observation rate  $p_{\text{obs}}$  so that each  $(i, j)$ ,  $i > j$  is included in  $\Omega$  independently with probability  $p_{\text{obs}}$ , and  $\Omega$  is assumed to be independent of  $\eta$ . While our goal is to accommodate general noise distributions, it is often helpful to bear in mind the following concrete example—termed a *random corruption model*—such that

$$y_{i,j} = \begin{cases} x_i - x_j \bmod m, & \text{with probability } \pi_0, \\ \text{Unif}(m), & \text{else,} \end{cases} \quad (3)$$

with  $\text{Unif}(m)$  being the uniform distribution over  $\{0, 1, \dots, m-1\}$ . The parameter  $\pi_0$  is called the *non-corruption rate*, since with probability  $1 - \pi_0$  the observation behaves like a random noise carrying no information whatsoever. The aim is to simultaneously recover all  $\{x_i\}$  based on the measurements  $\{y_{i,j}\}$  except for the global offset that is unrecoverable.<sup>3</sup>

Assuming the noise  $\{\eta_{i,j}\}$  has independent components, one can seek the MLE by solving

$$\begin{aligned} & \underset{\{z_i\}}{\text{maximize}} && \sum_{(i,j) \in \Omega} \ell(z_i, z_j; y_{i,j}) \\ & \text{subject to} && z_i \in \{1, \dots, m\}, \quad i = 1, \dots, n, \end{aligned} \quad (4)$$

where  $\ell(z_i, z_j; y_{i,j})$  represents the log-likelihood (or some other equivalent function) of the candidate solution  $(z_i, z_j)$  given the outcome  $y_{i,j}$ . For instance, under the random corruption model (3) one can write

$$\ell(z_i, z_j; y_{i,j}) = \begin{cases} \log\left(\pi_0 + \frac{1-\pi_0}{m}\right), & \text{if } z_i - z_j = y_{i,j} \bmod m, \\ \log\left(\frac{1-\pi_0}{m}\right), & \text{else.} \end{cases} \quad (5)$$

Throughout the rest of the paper, we set  $\ell(z_i, z_j; y_{i,j}) \equiv 0$  whenever  $(i, j) \notin \Omega$ .

This joint alignment problem finds applications in multiple domains. To begin with, the binary case (i.e.  $m = 2$ ) deserves special attention, as it reduces to a graph partitioning problem. For instance, in a

<sup>2</sup>We say  $\Omega$  is symmetric if  $(i, j) \in \Omega$  implies  $(j, i) \in \Omega$  for any  $i$  and  $j$ .

<sup>3</sup>Specifically, it is impossible to distinguish the  $m$  sets of inputs  $\{x_i\}_{1 \leq i \leq n}$ ,  $\{x_i - 1\}_{1 \leq i \leq n}$ ,  $\dots$ , and  $\{x_i - m + 1\}_{1 \leq i \leq n}$  given only pairwise differences.

community detection scenario in which one wishes to partition all users into two clusters, the variables  $\{x_i\}$  to recover indicate the cluster assignments for each user, while  $y_{i,j}$  represents the friendship between two users  $i$  and  $j$  (e.g. [24–39]). Another example is the problem of water-fat separation in medical imaging, which aims to partition all pixels of an MRI image into two parts representing water (denoted by  $x_i = 0$ ) and fat signals (denoted by  $x_i = 1$ ), respectively. The task takes as input some pre-computed pairwise cost functions  $-\ell(x_i, x_j)$ , which provides information about whether two pixels are of the same type or not; see [40–42] for details.

Moving beyond the binary case, this problem is motivated by the need of jointly aligning multiple images/shapes/pictures that arises in various fields. Imagine a sequence of images of the same physical instance (e.g. a building, a molecule), where each  $x_i$  represents the orientation of the camera when taking the  $i$ th image. A variety of computer vision tasks (e.g. 3D reconstruction from multiple scenes) or structural biology applications (e.g. cryo-electron microscopy) rely upon joint alignment of these images; or equivalently, joint recovery of the camera orientations associated with each image. Practically, it is often easier to estimate the relative camera orientation between a pair of images using raw features [43–45]. The problem then boils down to this: how to jointly aggregate such pairwise information in order to improve the collection of camera pose estimates?

### 1.3 Our contributions

In this work, we propose to compute the MLE (4) via a novel nonconvex procedure. Informally, the procedure starts by lifting each variable  $x_i \in [m]$  to higher dimensions such that distinct values are represented in orthogonal directions, and then encodes the log-likelihood function  $\ell(x_i, x_j; y_{i,j})$  for each  $(i, j)$  by an  $m \times m$  matrix. This way of representation allows to recast the MLE as a binary quadratic program or, equivalently, a constrained principal component analysis (PCA) problem. We then attempt optimization by means of projected power iterations, following an initial guess obtained via suitable low-rank factorization. This procedure proves effective in computing the MLE for a broad family of statistical models, and might be interesting for many other Boolean assignment problems beyond joint alignment.

## 2 Algorithm: projected power method

In this section, we present a nonconvex procedure to solve the MLE (4), which entails a series of projected power iterations over a higher-dimensional space. In what follows, this algorithm will be termed a *projected power method* (PPM).

### 2.1 Matrix representation of the MLE

The MLE admits an alternative matrix representation that is often more amenable to computation. To begin with, each state  $z_i \in \{1, \dots, m\}$  can be represented by a binary-valued vector  $\mathbf{z}_i \in \{0, 1\}^m$  such that

$$z_i = j \iff \mathbf{z}_i = \mathbf{e}_j \in \mathbb{R}^m, \quad (6)$$

where  $\mathbf{e}_1 = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$ ,  $\mathbf{e}_2 = \begin{bmatrix} 0 \\ 1 \\ \vdots \\ 0 \end{bmatrix}$ ,  $\dots$ ,  $\mathbf{e}_m = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix}$  are the standard basis vectors. In addition, for each

pair  $(i, j)$ , one can introduce an input matrix  $\mathbf{L}_{i,j} \in \mathbb{R}^{m \times m}$  to encode  $\ell(z_i, z_j; y_{i,j})$  given all possible input combinations of  $(z_i, z_j)$ ; that is,

$$(\mathbf{L}_{i,j})_{\alpha,\beta} := \ell(z_i = \alpha, z_j = \beta; y_{i,j}), \quad 1 \leq \alpha, \beta \leq m. \quad (7)$$

For example, when specialized to the random corruption model with binary alphabet ( $m = 2$ ),

$$\mathbf{L}_{i,j} = \begin{bmatrix} \log \frac{1+\pi_0}{2} & \log \frac{1-\pi_0}{2} \\ \log \frac{1-\pi_0}{2} & \log \frac{1+\pi_0}{2} \end{bmatrix} \text{ if } y_{i,j} = 0; \quad \text{and} \quad \mathbf{L}_{i,j} = \begin{bmatrix} \log \frac{1-\pi_0}{2} & \log \frac{1+\pi_0}{2} \\ \log \frac{1+\pi_0}{2} & \log \frac{1-\pi_0}{2} \end{bmatrix} \text{ if } y_{i,j} = 1.$$

By convention, we take  $\mathbf{L}_{i,i} \equiv \mathbf{0}$  for all  $1 \leq i \leq n$  and  $\mathbf{L}_{i,j} \equiv \mathbf{0}$  for all  $(i,j) \notin \Omega$ .

The preceding notation enables the quadratic form representation

$$\ell(z_i, z_j; y_{i,j}) = \mathbf{z}_i^\top \mathbf{L}_{i,j} \mathbf{z}_j.$$

For notational simplicity, we stack all the  $\mathbf{z}_i$ 's and the  $\mathbf{L}_{i,j}$ 's into a concatenated vector and matrix

$$\mathbf{z} = \begin{bmatrix} \mathbf{z}_1 \\ \vdots \\ \mathbf{z}_n \end{bmatrix} \in \mathbb{R}^{nm} \quad \text{and} \quad \mathbf{L} = \begin{bmatrix} \mathbf{L}_{1,1} & \cdots & \mathbf{L}_{1,n} \\ \vdots & \ddots & \vdots \\ \mathbf{L}_{n,1} & \cdots & \mathbf{L}_{n,n} \end{bmatrix} \in \mathbb{R}^{nm \times nm} \quad (8)$$

respectively, representing the states and log-likelihoods altogether. As a consequence, the MLE can be succinctly recast as a binary quadratic program, namely,

$$\begin{aligned} & \text{maximize}_{\mathbf{z}} \quad \mathbf{z}^\top \mathbf{L} \mathbf{z} \\ & \text{subject to} \quad \mathbf{z}_i \in \{\mathbf{e}_1, \dots, \mathbf{e}_m\}, \quad i = 1, \dots, n. \end{aligned} \quad (9)$$

This representation is appealing due to the simplicity of the objective function regardless of the landscape of  $\ell(\cdot, \cdot)$ , which allows one to focus on quadratic optimization rather than optimizing the (possibly complicated) function  $\ell(\cdot, \cdot)$  directly.

There are other families of  $\mathbf{L}$  which also lead to the MLE. We single out a simple, yet, important family obtained by enforcing global scaling and offset of  $\mathbf{L}$ . Specifically, the solution to (9) remains unchanged if each  $\mathbf{L}_{i,j}$  is replaced by<sup>4</sup>

$$\mathbf{L}_{i,j} \leftarrow a\mathbf{L}_{i,j} + b \cdot \mathbf{1}\mathbf{1}^\top, \quad \forall (i,j) \in \Omega \quad (10)$$

for some numerical values  $a > 0$  and  $b \in \mathbb{R}$ . Take the random corruption model as an example: a feasible form within this family is given by

$$(\mathbf{L}_{i,j})_{\alpha,\beta} = \begin{cases} 1, & \text{if } \alpha - \beta = y_{i,j} \bmod m, \\ 0, & \text{else,} \end{cases} \quad \forall (i,j) \in \Omega; \quad (11)$$

in words,  $\mathbf{L}_{i,j}$  is a cyclic permutation matrix obtained by circularly shifting the identity matrix  $\mathbf{I}_m$  by  $y_{i,j}$  positions. This form (11) does not rely on the non-corruption parameter  $\pi_0$ , and is hence practically more appealing. Another important instance in this family is the debiased version of  $\mathbf{L}$ —denoted by  $\mathbf{L}^{\text{debias}}$ —defined as follows,

$$\mathbf{L}_{i,j}^{\text{debias}} = \mathbf{L}_{i,j} - \frac{\mathbf{1}^\top \mathbf{L}_{i,j} \mathbf{1}}{m^2} \cdot \mathbf{1}\mathbf{1}^\top, \quad 1 \leq i, j \leq n, \quad (12)$$

which essentially removes the empirical average of the log-likelihoods in each entry.

## 2.2 Algorithm

One can interpret the binary quadratic program (9) as finding the principal component of  $\mathbf{L}$  subject to certain structural constraints. This motivates us to tackle this constrained PCA problem by means of a power method, with the assistance of appropriate regularization to enforce the structural constraints. More precisely, we consider the following procedure, which starting from a suitable initialization  $\mathbf{z}^{(0)}$  follows the update rule

$$\mathbf{z}^{(t+1)} = \mathcal{P}_{\Delta^n} \left( \mu_t \mathbf{L} \mathbf{z}^{(t)} \right), \quad \forall t \geq 0 \quad (13)$$

for some scaling parameter  $\mu_t$ . Here,  $\mathcal{P}_{\Delta^n}$  represents block-wise projection onto the standard simplex, namely, for any vector  $\mathbf{z} = [\mathbf{z}_i]_{1 \leq i \leq n} \in \mathbb{R}^{nm}$ ,

$$\mathcal{P}_{\Delta^n}(\mathbf{z}) := \begin{bmatrix} \mathcal{P}_{\Delta}(\mathbf{z}_1) \\ \vdots \\ \mathcal{P}_{\Delta}(\mathbf{z}_n) \end{bmatrix}, \quad (14)$$

<sup>4</sup>This is because  $\mathbf{z}_i^\top (a\mathbf{L}_{i,j} + b\mathbf{1}\mathbf{1}^\top) \mathbf{z}_j = a\mathbf{z}_i^\top \mathbf{L}_{i,j} \mathbf{z}_j + b$  given that  $\mathbf{1}^\top \mathbf{z}_i = \mathbf{1}^\top \mathbf{z}_j = 1$ .

---

**Algorithm 1** Projected power method.

---

**Input:** the input matrix  $\mathbf{L} = [\mathbf{L}_{i,j}]_{1 \leq i,j \leq n}$ ; the scaling factors  $\{\mu_t\}_{t \geq 0}$ .

**Initialize**  $\mathbf{z}^{(0)}$  to be  $\mathcal{P}_{\Delta^n}(\mu_0 \hat{\mathbf{z}})$  as defined in (14), where  $\hat{\mathbf{z}}$  is a random column of the best rank- $m$  approximation  $\hat{\mathbf{L}}$  of  $\mathbf{L}$ .

**Loop:** for  $t = 0, 1, \dots, T-1$  do

$$\mathbf{z}^{(t+1)} = \mathcal{P}_{\Delta^n}(\mu_t \mathbf{L} \mathbf{z}^{(t)}), \quad (17)$$

where  $\mathcal{P}_{\Delta^n}(\cdot)$  is defined in (14).

**Output**  $\{\hat{x}_i\}_{1 \leq i \leq n}$ , where  $\hat{x}_i$  is the index of the largest entry of the block  $\mathbf{z}_i^{(T)}$ .

---

where  $\mathcal{P}_{\Delta}(\mathbf{z}_i)$  is the projection of  $\mathbf{z}_i \in \mathbb{R}^m$  onto the standard simplex

$$\Delta := \{\mathbf{s} \in \mathbb{R}^m \mid \mathbf{1}^\top \mathbf{s} = 1; \mathbf{s} \text{ is non-negative}\}. \quad (15)$$

In particular, when  $\mu_t \rightarrow \infty$ ,  $\mathcal{P}_{\Delta}(\mu_t \mathbf{z}_i)$  reduces to a *rounding procedure*, which sets the largest entry of  $\mathbf{z}_i$  to one and all others to zero.

The key advantage of the PPM is its computational efficiency: the most expensive step in each iteration lies in matrix multiplication, which can be completed in nearly linear time; i.e. in time<sup>5</sup>  $O(|\Omega|m \log m)$ . This arises from the fact that each block  $\mathbf{L}_{i,j}$  is circulant so that we can compute a matrix-vector product using at most two  $m$ -point FFTs. The projection step  $\mathcal{P}_{\Delta}$  can be performed in  $O(m \log m)$  flops via a sorting-based algorithm (e.g. [46, Figure 1]), and hence  $\mathcal{P}_{\Delta^n}(\cdot)$  is much cheaper than the matrix-vector multiplication  $\mathbf{L} \mathbf{z}^{(t)}$  given that  $|\Omega| \gg n$  occurs in most applications.

One important step towards guaranteeing rapid convergence is to identify a decent initial guess  $\mathbf{z}^{(0)}$ . This is accomplished by low-rank factorization as follows:

1. Compute the best rank- $m$  approximation of the input matrix  $\mathbf{L}$ , namely,

$$\hat{\mathbf{L}} := \arg \min_{\mathbf{M}: \text{rank}(\mathbf{M}) \leq m} \|\mathbf{M} - \mathbf{L}\|_{\text{F}}. \quad (16)$$

2. Pick a random column  $\hat{\mathbf{z}}$  of  $\hat{\mathbf{L}}$  and set the initial guess as  $\mathbf{z}^{(0)} = \mathcal{P}_{\Delta^n}(\mu_0 \hat{\mathbf{z}})$ .

**Remark 1.** Alternatively, one can take  $\hat{\mathbf{L}}$  to be the best rank- $(m-1)$  approximation of the debiased input matrix  $\mathbf{L}^{\text{debias}}$  defined in (12), which can be computed in a slightly faster manner.

The main motivation comes from the (approximate) low-rank structure of the input matrix. As we shall shortly see,  $\mathbb{E}[\mathbf{L}]$  is approximately of rank  $m$ . Therefore, a low-rank approximation of  $\mathbf{L}$  is expected to preserve much of the information in  $\mathbb{E}[\mathbf{L}]$ , provided that the noise is not exceedingly large.

The low-rank factorization step can be performed efficiently via the method of orthogonal iteration (also called block power method) [47, Section 7.3.2]. Each power iteration consists of a matrix product of the form  $\mathbf{L}\mathbf{U}$  as well as a QR decomposition of some matrix  $\mathbf{V}$ , where  $\mathbf{U}, \mathbf{V} \in \mathbb{R}^{nm \times m}$ . The matrix product can be computed in  $O(|\Omega|m^2 \log m)$  flops with the assistance of  $m$ -point FFTs, whereas the QR decomposition takes time  $O(nm^3)$ . In summary, each power iteration runs in time  $O(|\Omega|m^2 \log m + nm^3)$ . Consequently, the matrix product constitutes the main computational cost when  $m \lesssim |\Omega| \log m/n$ , while the QR decomposition becomes the bottleneck when  $m \gg |\Omega| \log m/n$ .

The whole algorithm is summarized in Algorithm 1. There is of course the question of what  $\{\mu_t\}$  to use, which we defer to Section 3.

We remark that many variants of the power method have been proposed to tackle other generalized eigenproblems, including but not limited to sparse PCA [48–50], the hidden clique problem [51], water-fat separation [41], phase synchronization [52]. These procedures typically combine power iterations with additional procedures to exploit sparsity or enforce other feasibility constraints. The current work adds a new instance to this growing family of nonconvex methods, which is provably effective for a class of discrete problems.

---

<sup>5</sup>Here and throughout, the standard notion  $f(n) = O(g(n))$  or  $f(n) \lesssim g(n)$  mean there exists a constant  $c$  such that  $|f(n)| \leq cg(n)$ ;  $f(n) = o(g(n))$  means  $\lim_{n \rightarrow \infty} f(n)/g(n) = 0$ ;  $f(n) \gtrsim g(n)$  means there exists a constant  $c$  such that  $|f(n)| \geq cg(n)$ ; and  $f(n) \asymp g(n)$  means there exist constants  $c_1, c_2 > 0$  such that  $c_1 g(n) \leq |f(n)| \leq c_2 g(n)$ .

### 3 Main results

This section is devoted to quantifying the recovery performance of the projected power method. The key finding is that not only is the PPM much more practical than computing the MLE directly, it is capable of achieving nearly identical statistical accuracy as the MLE in a variety of scenarios.

#### 3.1 Random corruption model

Our theoretical study starts with the simple random corruption model (3). Notably, this model somehow corresponds to the worst-case situation since the uniform noise enjoys the highest entropy among all distributions over a fixed range, thus forming a reasonable benchmark for practitioners.

Recall that the noise level in this model is dictated by a single parameter  $\pi_0$ . We will show that the PPM is guaranteed to work even when the non-corruption rate  $\pi_0$  is vanishingly small, which is the scenario where almost all acquired measurements behave like random noise. While Algorithm 1 can certainly be implemented using the log-likelihoods (7), we recommend taking (11) as the input matrix in this case, as (11) is parameter free and, hence, more practical.

To state our results, we find it convenient to introduce a block sparsity metric. Specifically, the *block sparsity* of a vector  $\mathbf{h} = \{\mathbf{h}_i\}_{1 \leq i \leq n}$  is defined and denoted by

$$\|\mathbf{h}\|_{*,0} := \sum_{i=1}^n \mathbb{I}\{\mathbf{h}_i \neq 0\},$$

where  $\mathbb{I}\{\cdot\}$  is the indicator function. Since one can only hope to recover  $\mathbf{x}$  up to some global offset, we define the *misclassification rate* as the normalized block sparsity of the estimation error modulo the global shift

$$\text{MCR}(\hat{\mathbf{x}}, \mathbf{x}) := \frac{1}{n} \min_{0 \leq l < m} \|\hat{\mathbf{x}} - \text{shift}_l(\mathbf{x})\|_{*,0}. \quad (18)$$

Here,  $\text{shift}_l(\mathbf{x}) := [\text{shift}_l(\mathbf{x}_i)]_{1 \leq i \leq n} \in \mathbb{R}^{mn}$ , where  $\text{shift}_l(\mathbf{x}_i) \in \mathbb{R}^m$  is obtained by circularly shifting the entries of  $\mathbf{x}_i \in \mathbb{R}^m$  by  $l$  positions. With these notations in place, our result is this:

**Theorem 1.** *Consider the random corruption model (3) and the input matrix  $\mathbf{L}$  given in (11). Fix  $m > 0$ , and suppose  $p_{\text{obs}} > c_1 \log n/n$  and  $\mu_t > c_2/\sigma_2(\mathbf{L})$  for some sufficiently large constants  $c_1, c_2 > 0$ . Then there exists some absolute constant  $0 < \rho < 1$  such that with probability approaching one as  $n$  scales, the iterates of Algorithm 1 obey*

$$\text{MCR}(\mathbf{z}^{(t)}, \mathbf{x}) \leq 0.49\rho^t, \quad \forall t \geq 0, \quad (19)$$

provided that the non-corruption rate  $\pi_0$  exceeds<sup>6</sup>

$$\pi_0 > 2\sqrt{\frac{1.01 \log n}{mnp_{\text{obs}}}}. \quad (20)$$

**Remark 2.** Here and throughout,  $\sigma_i(\mathbf{L})$  is the  $i$ th largest singular value of  $\mathbf{L}$ . In fact, one can often replace  $\sigma_2(\mathbf{L})$  with  $\sigma_i(\mathbf{L})$  for other  $2 \leq i < m$ . But  $\sigma_1(\mathbf{L})$  is usually not a good choice unless we employ the debiased version of  $\mathbf{L}$  instead, because  $\sigma_1(\mathbf{L})$  typically corresponds to the “direct current” component of  $\mathbf{L}$  which could be excessively large. In addition, we note that  $\sigma_i(\mathbf{L})$  ( $i \leq m$ ) have been computed during spectral initialization and, as a result, will not introduce extra computational cost.

**Remark 3.** The contraction rate  $\rho$  can actually be as small as  $O(1/(\pi_0^2 n p_{\text{obs}}))$ , which is at most  $O(1/\log n)$  if  $m$  is fixed and if the condition (20) holds.

According to Theorem 1, convergence to the ground truth can be expected in at most  $O(\log n)$  iterations which, together with the per-iteration cost (which is on the order of  $|\Omega|$  since  $\mathbf{L}_{i,j}$  is a cyclic permutation matrix), shows that the computational complexity of the iterative stage is at most  $O(|\Omega| \log n)$ . This is nearly optimal since even reading all the data and likelihood values take time about the order of  $|\Omega|$ . All of this happens as soon as the non-corruption rate exceeds  $1 - O\left(\sqrt{\frac{\log n}{mnp_{\text{obs}}}}\right)$ , uncovering the remarkable ability of the PPM to tolerate and correct dense input errors.

<sup>6</sup>Theorem 1 continues to hold if we replace 1.01 with any other constant in  $(1, \infty)$ .

As we shall see later in Section 6, Theorem 1 holds if the algorithm starts with any other initial guess  $\mathbf{z}^{(0)}$  obeying

$$\text{MCR}(\mathbf{z}^{(0)}, \mathbf{x}) \leq 0.49, \quad (21)$$

irrespective of whether  $\mathbf{z}^{(0)}$  is independent of the data  $\{y_{i,j}\}$  or not. Therefore, it often suffices to run the power method for a constant number of iterations during the initialization stage, which can be completed in  $O(|\Omega|)$  flops when  $m$  is fixed. The broader implication is that Algorithm 1 remains successful if one adopts other initialization that can enter the basin of attraction.

Finally, our result is sharp: to be sure, the error-correction capability of the projected power method is statistically optimal, as revealed by the following converse result.

**Theorem 2.** *Consider the random corruption model (3) with any fixed  $m > 0$ , and suppose  $p_{\text{obs}} > c_1 \log n/n$  for some sufficiently large constant  $c_1 > 0$ . If <sup>7</sup>*

$$\pi_0 < 2\sqrt{\frac{0.99 \log n}{mn p_{\text{obs}}}}, \quad (22)$$

*then the minimax probability of error*

$$\inf_{\hat{\mathbf{x}}} \max_{x_i \in [m], 1 \leq i \leq n} \mathbb{P}(\text{MCR}(\hat{\mathbf{x}}, \mathbf{x}) > 0 \mid \mathbf{x}) \rightarrow 1, \quad \text{as } n \rightarrow \infty,$$

*where the infimum is taken over all estimators and  $\mathbf{x}$  is the vector representation of  $\{x_i\}$  as before.*

As mentioned before, the binary case  $m = 2$  bears some similarity with the community detection problem in the presence of two communities. Arguably the most popular model for community detection is the stochastic block model (SBM), where any two vertices within the same cluster (resp. across different clusters) are connected by an edge with probability  $p$  (resp.  $q$ ). The asymptotic limits for both exact and partial recovery have been extensively studied [26, 28, 29, 33, 38, 53–56]. We note, however, that the primary focus of community detection lies in the sparse regime  $p, q \asymp 1/n$  or logarithmic sparse regime (i.e.  $p, q \asymp \log n/n$ ), which is contrast to the joint alignment problem in which the measurements are often considerably denser. There are, however, a few theoretical results that extend to the dense regime, e.g. [28]. To facilitate comparison, by setting  $p = \frac{1+\pi_0}{2}$  and  $q = \frac{1-\pi_0}{2}$ , the SBM reduces to the random corruption model with  $p_{\text{obs}} = 1$ . One can easily verify that the limit  $\pi_0 = \sqrt{2 \log n/n}$  we derive matches the recovery threshold given in<sup>8</sup> [28, Theorem 2.5 and Proposition 2.9].

## 3.2 More general models

Theorems 1-2 are special instances of a set of general results that allow to accommodate a far more general family of noise models, as we will present in this subsection. Unless otherwise noted, we restrict attention to the class of *symmetric* noise distributions obeying

$$P_0(y) = P_0(-y), \quad \forall y, \quad (23)$$

which largely simplifies the exposition.

### 3.2.1 Key metrics

The feasibility of accurate recovery necessarily depends on the noise distribution or, more precisely, the distinguishability of the output distributions  $\{y_{i,j}\}$  given distinct inputs. In particular, there are  $m$  distributions  $\{P_l\}_{0 \leq l < m}$  that we'd like to emphasize, where  $P_l(\cdot)$  represents the distribution of  $y_{i,j}$  conditional on

<sup>7</sup>Theorem 2 continues to hold if we replace 0.99 with any other constant in  $(0,1)$ .

<sup>8</sup>Note that the model studied in [28] is an SBM with  $2n$  vertices with  $n$  vertices belonging to each cluster. Therefore, the threshold characterization [28, Proposition 2.9] should read

$$\frac{\sigma \sqrt{n/2}}{p-q} \exp\left(-\frac{n(p-q)^2}{4\sigma^2}\right) \rightarrow 0$$

when applied to our setting, with  $\sigma := \sqrt{p(1-q) + q(1-p)} \approx \sqrt{1/2}$ .

$x_i - x_j = l$ . Alternatively,  $P_l$  is also the  $l$ -shifted distribution of the noise  $\eta_{i,j}$  given by

$$P_l(y) := \mathbb{P}(y_{i,j} = y \mid x_i - x_j = l) = \mathbb{P}(\eta_{i,j} = y - l). \quad (24)$$

Here and below, we write  $a - b$  and  $a - b \bmod m$  interchangeably whenever it is clear from the context, and adopt the cyclic notation  $c_l = c_{l+m}$  ( $l \in \mathbb{Z}$ ) for any quantity taking the form  $c_l$ .

We would like to quantify the distinguishability of these distributions via some distance metric. One candidate is the Kullback–Leibler (KL) divergence defined by

$$\text{KL}(P_i \parallel P_l) := \sum_y P_i(y) \log \frac{P_i(y)}{P_l(y)}, \quad 0 \leq i, l < m, \quad (25)$$

which plays an important role in our main theory.

### 3.2.2 Performance guarantees

We now proceed to the main findings. To simplify matters, we shall concern ourselves primarily with the kind of noise distributions obeying the following assumption.

**Assumption 1.**  $m \min_y P_0(y)$  is bounded away from 0.

In words, Assumption 1 ensures that the noise density is not exceedingly lower than the average density  $1/m$  at any point. The reason why we introduce this assumption is two-fold. To begin with, this enables us to preclude the case where the entries of  $\mathbf{L}$ —or equivalently, the log-likelihoods—are too wild. For instance, if  $P_0(y) = 0$  for some  $y$ , then  $\log P_0(y) = -\infty$ , resulting in computational instability. Another reason is to simplify the analysis and exposition slightly, making it easier for the readers. We note, however, that this assumption is not crucial and can be dropped by means of a slight modification of the algorithm, which will be detailed later.

Another assumption that we would like to introduce is more subtle:

**Assumption 2.**  $\text{KL}_{\max}/\text{KL}_{\min}$  is bounded, where

$$\text{KL}_{\min} := \min_{1 \leq l < m} \text{KL}(P_0 \parallel P_l) \quad \text{and} \quad \text{KL}_{\max} := \max_{1 \leq l < m} \text{KL}(P_0 \parallel P_l), \quad (26)$$

Roughly speaking, Assumption 2 states that the mutual distances of the  $m$  possible output distributions  $\{P_l\}_{1 \leq l \leq m}$  lie within a reasonable dynamic range, so that one cannot find a pair of them that are considerably more separated than other pairs. Alternatively, it is understood that the variation of the log-likelihood ratio, as we will show later, is often governed by the KL divergence between the two corresponding distributions. From this point of view, Assumption 2 tells us that there is no submatrix of  $\mathbf{L}$  that is significantly more volatile than the remaining parts, which often leads to enhanced stability when computing the power iteration.

With these assumptions in place, we are in position to state our main result. It is not hard to see that Theorem 1 is an immediate consequence of the following theorem.

**Theorem 3.** Fix  $m > 0$ , and assume  $p_{\text{obs}} > c_1 \log n/n$  and  $\mu_t > c_2/\sigma_2(\mathbf{L})$  for some sufficiently large constants  $c_1, c_2 > 0$ . Under Assumptions 1–2, there exist some absolute constants  $0 < \rho, \nu < 1$  such that with probability tending to one as  $n$  scales, the iterates of Algorithm 1 with the input matrix (12) obey

$$\text{MCR}(\mathbf{z}^{(t)}, \mathbf{x}) \leq \nu \rho^t, \quad \forall t \geq 0, \quad (27)$$

provided that<sup>9</sup>

$$\text{KL}_{\min} \geq \frac{4.01 \log n}{np_{\text{obs}}}. \quad (28)$$

---

<sup>9</sup>Theorem 3 remains valid if we replace 4.01 by any other constant in  $(4, \infty)$ .



**Remark 4.** Alternatively, Theorem 3 can be stated in terms of other divergence metrics like the squared Hellinger distance  $H^2(\cdot, \cdot)$ . Specifically, Theorem 3 holds if the minimum squared Hellinger distance obeys

$$H_{\min}^2 := \min_{1 \leq l < m} H^2(P_0, P_l) > \frac{1.01 \log n}{np_{\text{obs}}}, \quad (29)$$

where  $H^2(P, Q) := \frac{1}{2} \sum_y (\sqrt{P(y)} - \sqrt{Q(y)})^2$ . We will show later in Lemma 3 that  $\text{KL}_{\min} \approx 4H_{\min}^2$ , which justifies the equivalence between (28) and (29).

The recovery condition (28) is non-asymptotic, and takes the form of a minimum KL divergence criterion. This is consistent with the understanding that the hardness of exact recovery often arises in differentiating minimally separated output distributions. Within at most  $O(\log n)$  projected power iterations, the PPM returns an estimate with absolutely no error, as soon as the minimum KL divergence exceeds some threshold. This threshold can be remarkably small when  $p_{\text{obs}}$  is large or, equivalently, when we have many pairwise measurements available.

Theorem 3 accommodates a broad class of noise models. Here we highlight a few examples to illustrate its generality. To begin with, it is self-evident that the random corruption model belongs to this class with  $\text{KL}_{\max}/\text{KL}_{\min} = 1$ . Beyond this simple model, we list two important families which satisfy Assumption 2 and which receive broad practical interest. This list, however, is by no means exhaustive.

- (1) A class of distributions that obey

$$\text{KL}_{\min} = \text{KL}(P_0 \parallel P_1) \quad \text{or} \quad \text{KL}_{\min} = \text{KL}(P_0 \parallel P_{-1}). \quad (30)$$

This says that the output distributions are the closest when the two corresponding inputs are minimally separated.

- (2) A class of *unimodal* distributions that satisfy

$$P_0(0) \geq P_0(1) \geq \dots \geq P_0(\lfloor m/2 \rfloor). \quad (31)$$

This says that the likelihood decays as the distance to the truth increases.

**Lemma 1.** Fix  $m > 0$ , and suppose Assumption 1 holds. Then the noise distribution satisfying either (30) or (31) obeys Assumption 2.

*Proof.* See Appendix B. □

### 3.2.3 Why Algorithm 1 works?

We pause here to gain some insights about Algorithm 1 and, in particular, why the minimum KL divergence has emerged as a key metric. Without loss of generality, we assume  $x_1 = \dots = x_n = 1$  to simplify the presentation.

Recall that Algorithm 1 attempts to find the constrained principal component of  $\mathbf{L}$ . To enable successful recovery, one would naturally hope the structure of the data matrix  $\mathbf{L}$  to reveal much information about the truth. In the limit of large samples, it is helpful to start by looking at the mean of  $\mathbf{L}_{i,j}$ , which is given by

$$\begin{aligned} \mathbb{E}[\mathbf{L}_{i,j}]_{\alpha,\beta} &= p_{\text{obs}} \mathbb{E}_{y \sim P_0} [\log P_{\alpha-\beta}(y)] = p_{\text{obs}} \mathbb{E}_{y \sim P_0} \left[ \log \frac{P_{\alpha-\beta}(y)}{P_0(y)} \right] + p_{\text{obs}} \mathbb{E}_{y \sim P_0} [\log P_0(y)] \\ &= p_{\text{obs}} [-\text{KL}_{\alpha-\beta} - \mathcal{H}(P_0)] \end{aligned} \quad (32)$$

for any  $i \neq j$  and  $1 \leq \alpha, \beta \leq m$ ; here and throughout,  $\mathcal{H}(P_0) := -\sum_y P_0(y) \log P_0(y)$  is the entropy functional, and

$$\text{KL}_l := \text{KL}(P_0 \parallel P_l), \quad 0 \leq l < m. \quad (33)$$

We can thus write

$$\mathbb{E}[\mathbf{L}] = p_{\text{obs}} \begin{bmatrix} \mathbf{0} & \mathbf{K} & \cdots & \mathbf{K} \\ \mathbf{K} & \mathbf{0} & \ddots & \vdots \\ \vdots & \ddots & \mathbf{0} & \mathbf{K} \\ \mathbf{K} & \cdots & \mathbf{K} & \mathbf{0} \end{bmatrix} \quad (34)$$

with  $\mathbf{K} \in \mathbb{R}^{m \times m}$  denoting a circulant matrix

$$\mathbf{K} := \underbrace{\begin{bmatrix} -\text{KL}_0 & -\text{KL}_{m-1} & \cdots & -\text{KL}_1 \\ -\text{KL}_1 & -\text{KL}_0 & \ddots & -\text{KL}_2 \\ \vdots & \ddots & \ddots & \vdots \\ -\text{KL}_{m-1} & \cdots & -\text{KL}_1 & -\text{KL}_0 \end{bmatrix}}_{:=\mathbf{K}^0} - \mathcal{H}(P_0) \mathbf{1} \cdot \mathbf{1}^\top. \quad (35)$$

It is easy to see that the largest entries of  $\mathbf{K}$  lie on the main diagonal, due to the fact that

$$-\text{KL}_0 = 0 \quad \text{and} \quad -\text{KL}_l = -\text{KL}(P_0 \parallel P_l) < 0 \quad (1 \leq l < m).$$

Consequently, for any column of  $\mathbb{E}[\mathbf{L}]$ , knowledge of the largest entries in each block reveals the relative positions across all  $\{x_i\}$ . Take the 2nd column of  $\mathbb{E}[\mathbf{L}]$  for example: all but the first blocks of this column attain the maximum values in their 2nd entries, telling us that  $x_2 = \cdots = x_n$ .

Given the noisy nature of the acquired data, one would further need to ensure that the true structure stands out from the noise. This hinges upon understanding when  $\mathbf{L}$  can serve as a reasonably good proxy for  $\mathbb{E}[\mathbf{L}]$  in the (projected) power iterations. Since we are interested in identifying the largest entries, the signal contained in each block—which is essentially the mean separation between the largest and second largest entries—is of size

$$p_{\text{obs}} \min_{1 \leq l < m} \mathbb{E}_{y \sim P_0} [\log P_0(y) - \log P_l(y)] = p_{\text{obs}} \min_{1 \leq l < m} \text{KL}_l = p_{\text{obs}} \text{KL}_{\min}.$$

The total signal strength is thus given by  $np_{\text{obs}} \text{KL}_{\min}$ . In addition, the variance in each measurement is bounded by

$$\max_{1 \leq l < m} \text{Var}_{y \sim P_0} [\log P_0(y) - \log P_l(y)] \stackrel{(a)}{\lesssim} \max_{1 \leq l < m} \text{KL}_l = \text{KL}_{\max},$$

where the inequality (a) will be demonstrated later in Lemma 3. From the semicircle law, the perturbation can be controlled by

$$\|\mathbf{L} - \mathbb{E}[\mathbf{L}]\| = O\left(\sqrt{np_{\text{obs}} \text{KL}_{\max}}\right), \quad (36)$$

which cannot exceed the size of the signal, namely,

$$np_{\text{obs}} \text{KL}_{\min} \gtrsim \sqrt{np_{\text{obs}} \text{KL}_{\max}}.$$

This condition reduces to

$$\text{KL}_{\min} \gtrsim \frac{1}{np_{\text{obs}}}$$

under Assumption 2, which is consistent with Theorem 3 up to some logarithmic factor.

### 3.2.4 Optimality

The preceding performance guarantee turns out to be information theoretically optimal in the asymptotic regime. In fact, the KL divergence threshold given in Theorem 3 is arbitrarily close to the information limit, a level below which every procedure is bound to fail in a minimax sense. We formalize this finding as a converse result:

**Theorem 4.** Fix  $m > 0$ . Let  $\{P_{l,n} : 0 \leq l < m\}_{n \geq 1}$  be a sequence of probability measures supported on a finite set  $\mathcal{Y}$ , where  $\inf_{y \in \mathcal{Y}, n, p} P_{l,n}(y)$  is bounded away from 0. Suppose that there exists  $1 \leq l < m$  such that

$$\text{KL}(P_{0,n} \parallel P_{l,n}) = \min_{1 \leq j < m} \text{KL}(P_{0,n} \parallel P_{j,n}) := \text{KL}_{\min,n}$$

for all sufficiently large  $n$ , and that  $p_{\text{obs}} > \frac{c_0 \log n}{n}$  for some sufficiently large constant  $c_0 > 0$ . If <sup>10</sup>

$$\text{KL}_{\min,n} < \frac{3.99 \log n}{np_{\text{obs}}}, \quad (37)$$

then the minimax probability of error

$$\inf_{\hat{\mathbf{x}}} \max_{x_i \in [m], 1 \leq i \leq n} \mathbb{P}(\text{MCR}(\hat{\mathbf{x}}, \mathbf{x}) > 0 \mid \mathbf{x}) \rightarrow 1, \quad \text{as } n \rightarrow \infty, \quad (38)$$

where the infimum is over all possible estimators and  $\mathbf{x}$  is the vector representation of  $\{x_i\}$  as usual.

### 3.3 Extension: removing Assumption 1

We return to Assumption 1. As mentioned before, an exceedingly small  $P_0(y)$  might result in unstable log-likelihoods, which suggests we regularize the data before running the algorithm. To this end, one alternative is to introduce a little more entropy to the samples so as to regularize the noise density, namely, we add a small level of random noise to yield

$$\tilde{y}_{i,j} \stackrel{\text{ind.}}{=} \begin{cases} y_{i,j}, & \text{with probability } 1 - \xi, \\ \text{Unif}(m), & \text{else,} \end{cases} \quad (39)$$

for some appropriate small constant  $\xi > 0$ . The distribution of the new data  $\tilde{y}_{i,j}$  given  $x_i - x_j = l$  is thus given by

$$\tilde{P}_l \leftarrow (1 - \xi)P_l + \xi \text{Unif}(m), \quad (40)$$

which effectively bumps  $\min P_0(y)$  up to  $(1 - \xi) \min P_0(y) + \xi/m$ . We then propose to run Algorithm 1 using the new data  $\{\tilde{y}_{i,j}\}$  and  $\{\tilde{P}_l\}$ , leading to the following performance guarantee.

**Theorem 5.** Take  $\xi > 0$  to be some sufficiently small constant, and suppose that Algorithm 1 operates upon  $\{\tilde{y}_{i,j}\}$  and  $\{\tilde{P}_l\}$ . Then Theorem 3 holds without Assumption 1.

*Proof.* See Appendix A. □

### 3.4 Extension: large- $m$ case

So far our study has focused on the case where the alphabet size  $m$  does not scale with  $n$ . There are, however, no shortage of situations where  $m$  is so large that it cannot be treated as a fixed constant. The encouraging news is that Algorithm 1 appears surprisingly competitive for the large- $m$  case as well. Once again, we begin with the random corruption model, and our analysis developed for fixed  $m$  immediately applies here.

**Theorem 6.** Suppose that  $m \gtrsim \log n$ ,  $m = O(\text{poly}(n))$ , and  $p_{\text{obs}} \geq c_1 \log^2 n / n$  for some sufficiently large constant  $c_1 > 0$ . Then Theorem 1 continues to hold with probability at least 0.99, as long as (20) is replaced by

$$\pi_0 > \frac{c_3}{\sqrt{np_{\text{obs}}}} \quad (41)$$

for some universal constant  $c_3 > 0$ . Here, 0.99 is arbitrary and can be replaced by any constant in  $(0, 1)$ .

The main message of Theorem 6 is that the error correction capability of the proposed method improves as the number  $n$  of unknowns grows. The quantitative bound (41) implies successful recovery even when an overwhelming fraction of the measurements are corrupted. Notably, when  $m$  is exceedingly large, Theorem 6 might shed light on the continuous joint alignment problem. In particular, there are two cases worth emphasizing:

---

<sup>10</sup>Theorem 4 continues to hold if we replace 3.99 with any other constant in  $(0, 4)$ .

- When  $m \gg n$  (e.g.  $m \gtrsim n^{10}$ ), the random corruption model converges to the following continuous spike model as  $n$  scales:

$$x_i \in [0, 1), \quad 1 \leq i \leq n \quad \text{and} \quad y_{i,j} = \begin{cases} x_i - x_j \bmod 1, & \text{with prob. } \pi_0, \\ \text{Unif}(0, 1), & \text{else,} \end{cases} \quad (i, j) \in \Omega. \quad (42)$$

This coincides with the setting studied in [44, 57] over the orthogonal group  $\text{SO}(2)$ , under the name of *synchronization* [37, 52, 58–60]. It has first been shown that the leading eigenvector of a certain Hermitian data matrix becomes positively correlated with the ground truth as long as  $\pi_0 > 1/\sqrt{n}$  when  $p_{\text{obs}} = 1$  [57]. In addition, a generalized power method—which is equivalent to projected gradient descent—provably converges to the solution of the nonconvex least-squares estimation, as long as the size of the noise is below some threshold [52, 59]. When it comes to exact recovery, Wang et al. prove that semidefinite relaxation succeeds as long as  $\pi_0 > 0.457$  [44, Theorem 4.1], a constant threshold irrespective of  $n$ . In contrast, the exact recovery performance of our approach—which operates over a lifted discrete space rather than  $\text{SO}(2)$ —improves with  $n$ , allowing  $\pi_0$  to be arbitrarily small when  $n$  is sufficiently large. On the other hand, the model (42) is reminiscent of the more general robust PCA problem [61, 62], which consists in recovering a low-rank matrix when a fraction of observed entries are corrupted. We have learned from the literature [63, 64] that perfect reconstruction is feasible and tractable even though a dominant portion of the observed entries may suffer from random corruption, which is consistent with our finding in Theorem 6.

- In the preceding spike model, the probability density of each measurement experiences an impulse around the truth. In a variety of realistic scenarios, however, the noise density might be more smooth rather than being spiky. This smoothness condition can be modeled by enforcing  $P_0(z) \asymp P_0(z + 1)$  for all  $z$ , so as to rule out any sharp jump. To satisfy this condition, we can at most take  $m \asymp \sqrt{np_{\text{obs}}}$  in view of Theorem 6. In some sense, this uncovers the “resolution” of our estimator under the smooth noise model: if we constrain the input domain to be the unit interval by letting  $x_i = 1, \dots, m$  represent the grid points  $0, \frac{1}{m}, \dots, \frac{m-1}{m}$ , respectively, then the PPM can recover each variable up to a resolution of

$$\frac{1}{m} \asymp \frac{1}{\sqrt{np_{\text{obs}}}}. \quad (43)$$

Notably, the discrete random corruption model has been investigated in prior literature [65, 66], with the best theoretical support derived for convex programming. Specifically, it has been shown by [66] that convex relaxation is guaranteed to work as soon as  $\pi_0 \gtrsim \frac{\log^2 n}{\sqrt{np_{\text{obs}}}}$ , which loses some logarithmic factor in comparison to our theory (41) developed for the projected power method.

Furthermore, Theorem 6 is an immediate consequence of a more general result:

**Theorem 7.** *Assume that  $m \gtrsim \log n$ ,  $m = O(\text{poly}(n))$ , and  $p_{\text{obs}} > c_1 \log^5 n/n$  for some sufficiently large constant  $c_1 > 0$ . Suppose that  $\mathbf{L}$  is replaced by  $\mathbf{L}^{\text{debias}}$  when computing the initial guess  $\mathbf{z}^{(0)}$  and  $\mu \geq c_3 \sqrt{m}/\sigma_m(\mathbf{L}^{\text{debias}})$  for some sufficiently large constant  $c_3 > 0$ . Then Theorem 3 continues to hold with probability exceeding 0.99, provided that (28) is replaced by*

$$\frac{\text{KL}_{\min}^2}{\max_{1 \leq l < m} \left\| \log \frac{P_0}{P_l} \right\|_1^2} \geq \frac{c_3}{np_{\text{obs}}} \quad (44)$$

for some universal constant  $c_3 > 0$ . Here,  $\left\| \log \frac{P_0}{P_l} \right\|_1 := \sum_y \left| \log \frac{P_0(y)}{P_l(y)} \right|$ , and 0.99 can be replaced by any constant in  $(0, 1)$ .

Some brief interpretations of (44) are in order. As discussed before, the quantity  $\text{KL}_{\min}$  represents the size of the signal. In contrast, the term  $\max_{1 \leq l < m} \left\| \log \frac{P_0}{P_l} \right\|_1^2$  controls the variability of each block of the data matrix  $\mathbf{L}_{i,j}$ ; to be more precise,

$$\mathbb{E} \left[ \left\| \mathbf{L}_{i,j} - \mathbb{E}[\mathbf{L}_{i,j}] \right\|^2 \right] \lesssim \max_{1 \leq l < m} \left\| \log \frac{P_0}{P_l} \right\|_1^2$$

as we will demonstrate in the proof. Thus, the left-hand side of (44) can be regarded as the signal-to-noise-ratio (SNR) experienced in each block  $\mathbf{L}_{i,j}$ . The recovery criterion is thus in terms of a lower threshold on the SNR, which can be vanishingly small in the regime considered in Theorem 7. We note that the general alignment problem has been studied in [45] as well, although the focus therein is to show the stability of semidefinite relaxation in the presence of random vertex noise.

We caution, however, that the performance guarantees presented in this subsection are in general not information-theoretically optimal. For instance, it has been shown [67] that the MLE succeeds as long as

$$\pi_0 \gtrsim \sqrt{\frac{\log n}{mnp_{\text{obs}}}}$$

in the regime where  $m \ll np_{\text{obs}}/\log n$ ; that is, the performance of the MLE improves as  $m$  increases. It is noteworthy that none of the polynomial algorithms proposed in prior works achieves optimal scaling in  $m$ . It remains to be seen whether this arises due to some drawback of the algorithms, or due to the existence of some inherent information-computation gap.

## 4 Numerical experiments

This section examines the empirical performance of the projected power method on both synthetic instances and real image data.

### 4.1 Synthetic experiments

To begin with, we conduct a series of Monte Carlo trials for various problem sizes under the random corruption model (3). Specifically, we vary the number  $n$  of unknowns, the input corruption rate  $1 - \pi_0$ , and the alphabet size  $m$ , with the observation rate set to be  $p_{\text{obs}} = 1$  throughout. For each  $(n, \pi_0, m)$  tuple, 20 Monte Carlo trials are conducted. In each trial, we draw each  $x_i$  uniformly at random over  $[m]$ , generate a set of measurements  $\{y_{i,j}\}_{1 \leq i, j \leq n}$  according to (3), and record the misclassification rate  $\text{MCR}(\hat{\mathbf{x}}, \mathbf{x})$  of Algorithm 1. The mean empirical misclassification rate is then calculated by averaging over 20 Monte Carlo trials.

Fig. 1 depicts the mean empirical misclassification rate when  $m = 2, 10, 20$ , and accounts for two choices of the scaling factors: (1)  $\mu_t \equiv 10/\sigma_2(\mathbf{L})$ , and (2)  $\mu_t \equiv \infty$ . In particular, the solid lines locate the asymptotic phase transitions for exact recovery predicted by our theory. In all cases, the empirical phase transition curves come closer to the analytical prediction as the problem size  $n$  increases.

Another noise model that we have studied numerically is a modified Gaussian model. Specifically, we set  $m = 5, 9, 15$ , and the random noise  $\eta_{i,j}$  is generated in such a way that

$$\mathbb{P}\{\eta_{i,j} = z\} \propto \exp\left(-\frac{z^2}{2\sigma^2}\right), \quad -\frac{m-1}{2} \leq z \leq \frac{m-1}{2}, \quad (45)$$

where  $\sigma$  controls the flatness of the noise density. We vary the parameters  $(\sigma, n, m)$ , take  $p_{\text{obs}} = 1$ , and experiment on two choices of scaling factors  $\mu_t \equiv 20/\sigma_m(\mathbf{L})$  and  $\mu_t \equiv \infty$ . The mean misclassification rate of the PPM is reported in Fig. 2, where the empirical phase transition matches the theory very well.

### 4.2 Joint shape alignment

Next, we return to the motivating application—i.e. joint image/shape alignment—of this work, and validate the applicability of the PPM on two datasets drawn from the ShapeNet repository [68]: (a) the Chair dataset (03001627), and (b) the Plane dataset (02691156). Specifically,  $n = 50$  shapes are taken from each dataset, and we randomly sample 8192 points from each shape as input features. Each shape is rotated in the  $x$ - $z$  plane by a random continuous angle  $\theta_i \in [0, 360^\circ)$ . Since the shapes in these datasets have high quality and low noise, we perturb the shape data by adding independent Gaussian noise  $\mathcal{N}(0, 0.2^2)$  to each coordinate of each point and use the perturbed data as inputs. This makes the task more challenging; for instance, the resulting SNR on the Chair dataset is around 0.945 (since the mean square values of each coordinate of the samples is 0.0378).

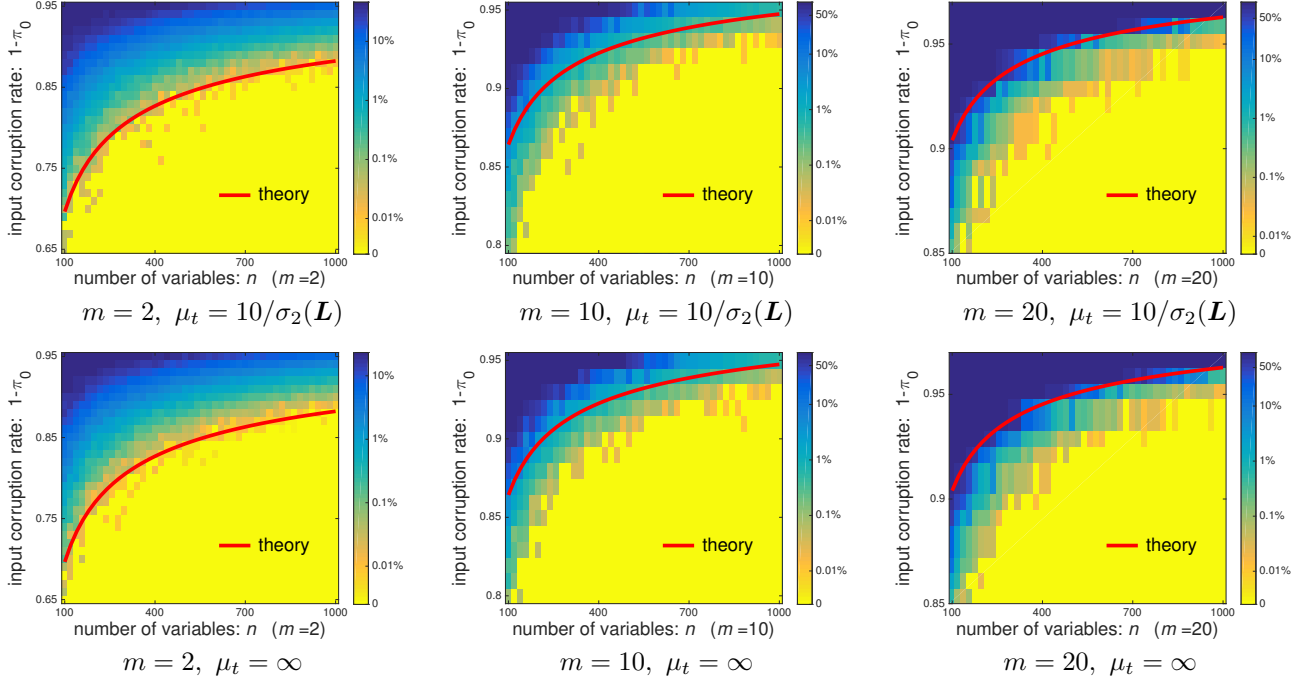


Figure 1: The empirical mean misclassification rate of Algorithm 1 under the random corruption model.

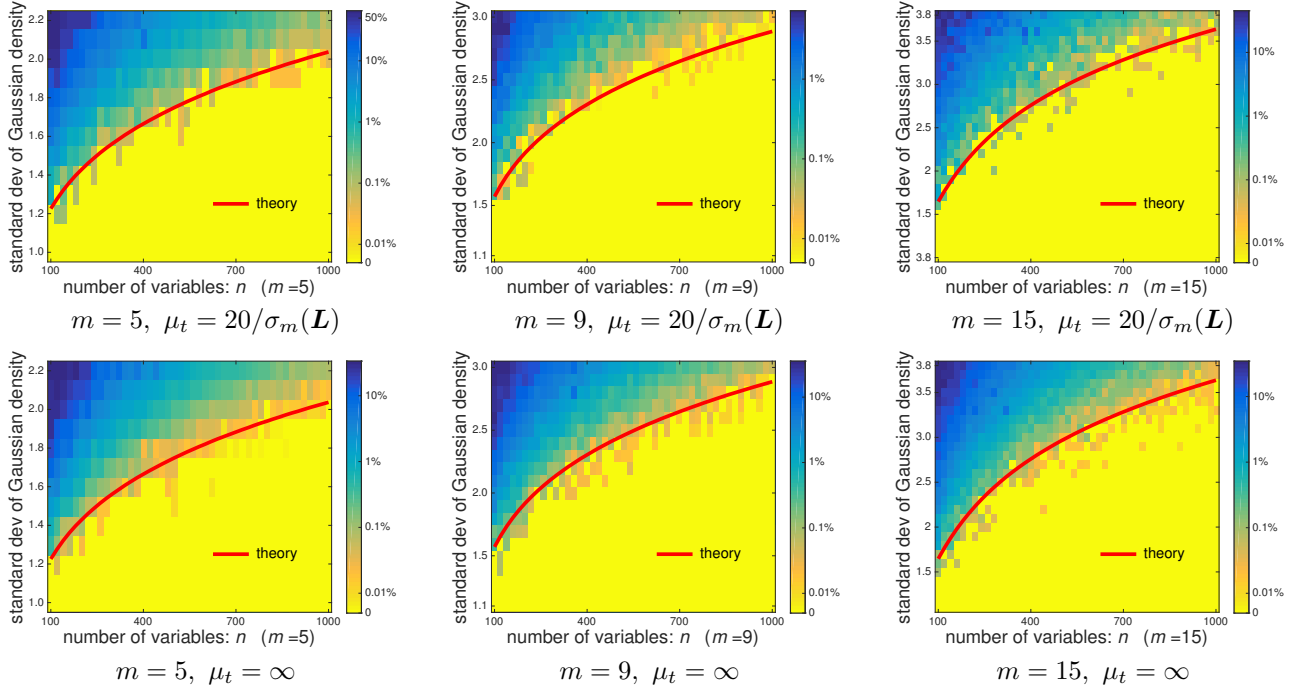


Figure 2: The empirical mean misclassification rate of Algorithm 1 under the modified Gaussian model.

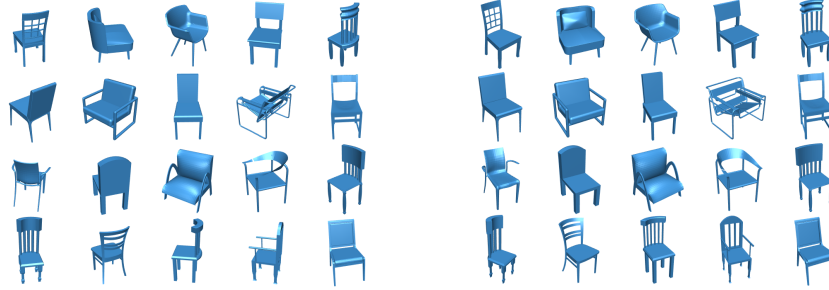


Figure 3: The performance of the PPM on a Chair dataset of  $n = 50$  shapes: (left) the first 20 input shapes; (right) the first 20 shapes after alignment.

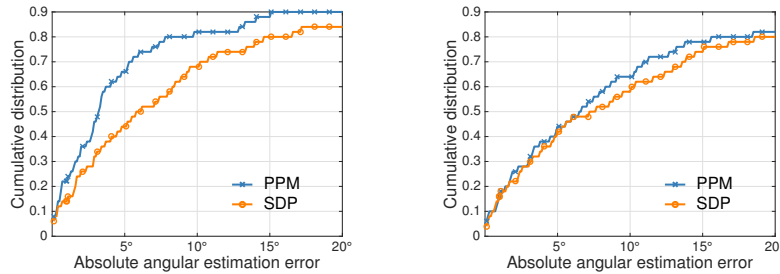


Figure 4: The cumulative distributions of the absolute angular estimation errors on: (left) the Plane dataset, and (right) the Chair dataset.

To apply the projected power method, we discretize the angular domain  $[0, 360^\circ)$  by  $m = 32$  points, so that  $x_i = j$  ( $1 \leq j \leq 32$ ) represents an angle  $\theta_i = j360^\circ/32$ . Following the procedure adopted in<sup>11</sup> [43], we compute the pairwise cost (i.e.  $-\ell(z_i, z_j)$ ) using some nearest-neighbor distance metric; to be precise, we set  $-\ell(z_i, z_j)$  as the average nearest-neighbor squared distance between the samples of the  $i$ th and  $j$ th shapes, after they are rotated by  $\frac{z_i}{32}360^\circ$  and  $\frac{z_j}{32}360^\circ$ , respectively. Such pairwise cost functions have been widely used in computer graphics and vision, and one can regard it as assuming that the average nearest-neighbor distance follows some Gaussian distribution. Careful readers might remark that we have not specified  $\{y_{i,j}\}$  in this experiment. Practically, oftentimes we only have access to some pairwise potential/cost functions rather than  $\{y_{i,j}\}$ . Fortunately, all we need to run the algorithm is  $\ell(z_i, z_j)$ , or some proxy of  $\ell(z_i, z_j)$ .

Fig. 3 shows the first 20 representative shapes before and after joint alignment in the Chair dataset. As one can see, the shapes are aligned in a reasonably good manner. More quantitatively, Fig. 4 displays the cumulative distributions of the absolute angular estimation errors for both datasets. We have also reported in Fig. 4 the performance of semidefinite programming (SDP)—that is, the MatchLift algorithm presented in [66]. Note that the angular errors are measured as the distance to the un-discretized angles  $\{\theta_i\}$ , and are hence somewhat continuous. We see that for the PPM, 70% (resp. 44%) of the estimates on the Plane (resp. Chair) dataset have an error of  $5.5^\circ$  or lower, while the proportion is 48% (resp. 44%) for the SDP formulation. Recall that the resolution of the discretization is  $360^\circ/32 \approx 11^\circ$ , which would mean that all estimates with an error less than  $5.5^\circ$  are, in some sense, perfect recoveries.

Computationally, it takes around 2.4 seconds to run the PPM, while SDP (implemented using the alternating direction method of multipliers (ADMM)) runs in 895.6 seconds. All experiments are carried out on a MacBook Pro equipped with a 2.9 GHz Intel Core i5 and 8GB of memory.

### 4.3 Joint graph matching

The PPM is applicable to other combinatorial problems beyond joint alignment. We present here an example called joint graph matching [65,66,69–71]. Consider a collection of  $n$  images each containing  $m$  feature points,

<sup>11</sup>[https://github.com/huangqx/map\\_synchronization/](https://github.com/huangqx/map_synchronization/)

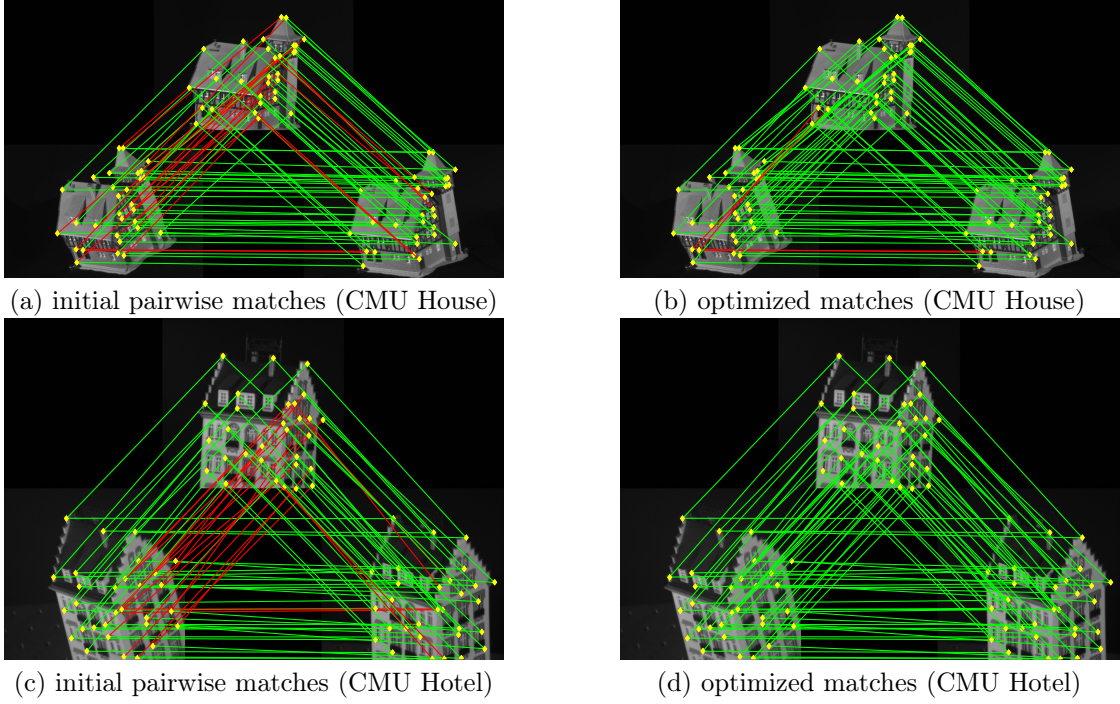


Figure 5: Comparisons between the input matches and the outputs of the PPM on the CMU House and Hotel datasets, with 3 representative images shown for each dataset. The yellow dots refer to the manually labeled feature points, while the green (resp. red) lines represent the set of matches consistent (resp. inconsistent) with the ground truth.

and suppose that there exists a one-to-one correspondence between the feature points in any pair of images. Many off-the-shelf algorithms are able to compute feature correspondence over the points in two images, and the joint matching problem concerns the recovery of a collection of globally consistent feature matches given these noisy pairwise matches. To put it mathematically, one can think of the ground truth as  $n$  permutation matrices  $\{\mathbf{X}_i \in \mathbb{R}^{m \times m}\}_{1 \leq i \leq n}$  each representing the feature mapping between an image and a reference, and the true feature correspondence over the  $i$ th and  $j$ th images can be represented by  $\mathbf{X}_i \mathbf{X}_j^\top$ . The provided pairwise matches between the features of two images are encoded by  $\mathbf{L}_{i,j} \in \mathbb{R}^{m \times m}$ , which is a noisy version of  $\mathbf{X}_i \mathbf{X}_j^\top$ . The goal is then to recover  $\{\mathbf{X}_i\}$ —up to some global permutation—given a set of pairwise observations  $\{\mathbf{L}_{i,j}\}$ . See [65] for more details.

This problem differs from joint alignment in that the ground truth  $\mathbf{X}_i$  is an  $m \times m$  permutation matrix. In light of this, we make two modifications to the algorithm: (i) we maintain the iterates  $\mathbf{Z}^t = [\mathbf{Z}_i^t]_{1 \leq i \leq n}$  as  $nm \times m$  matrices and replace  $\mathcal{P}_\Delta$  by  $\mathcal{P}_\Pi(\cdot)$  that projects each  $\mathbf{Z}_i^t \in \mathbb{R}^{m \times m}$  to the set of permutation matrices (via the Jonker-Volgenant algorithm [72]), which corresponds to hard rounding (i.e.  $\mu_t = \infty$ ) in power iterations (ii) the initial guess  $\mathbf{Z}^{(0)} \in \mathbb{R}^{nm \times m}$  is taken to be the projection of a random column block of  $\hat{\mathbf{L}}$  (which is the rank- $m$  approximation of  $\mathbf{L}$ ).

We first apply the PPM on two benchmark image datasets: (1) the CMU House dataset<sup>12</sup> consisting of  $n = 111$  images of a house, and (2) the CMU Hotel dataset<sup>13</sup> consisting of  $n = 101$  images of a hotel. Each image contains  $m = 30$  feature points that have been labeled consistently across all images. The initial pairwise matches, which are obtained through the Jonker-Volgenant algorithm, have mismatching rates of 13.36% (resp. 12.94%) for the House (resp. Hotel) dataset. Our algorithm allows to lower the mismatching rate to 3.25% (resp. 4.81%) for House (resp. Hotel). Some representative results from each dataset are depicted in Fig. 5.

Next, we turn to three shape datasets: (1) the Hand dataset containing  $n = 20$  shapes, (2) the Fourleg

<sup>12</sup><http://vasc.ri.cmu.edu/idb/html/motion/house/>

<sup>13</sup><http://vasc.ri.cmu.edu/idb/html/motion/hotel/index.html>



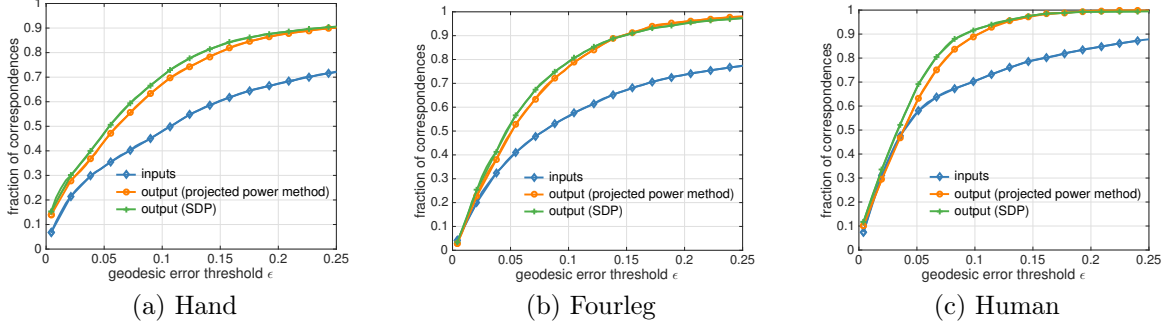


Figure 6: The fraction of correspondences whose normalized geodesic errors are smaller than some threshold  $\epsilon$ .

	Hand	Fourleg	Human
SDP	455.5 sec	1389.6 sec	368.9 sec
PPM	35.1 sec	76.8 sec	40.8 sec

Table 1: Runtime of SDP (implemented using ADMM) and the PPM on 3 benchmark datasets, when carried out on a MacBook Pro equipped with a 2.9 GHz i5 and 8GB of memory.

dataset containing  $n = 20$  shapes, and (3) the Human dataset containing  $n = 18$  shapes, all of which are drawn from the collection SHREC07 [73]. We set  $m = 64$ ,  $m = 96$ , and  $m = 64$  feature points for Hand, Fourleg and Human, respectively, and follow the shape sampling and pairwise matching procedures described in [65]<sup>14</sup>. To evaluate the matching performance, we report the fraction of output matches whose normalized geodesic errors (see [65, 69]) are below some threshold  $\epsilon$ , with  $\epsilon$  ranging from 0 to 0.25. For the sake of comparisons, we plot in Fig. 6 the quality of the initial matches, the matches returned by the projected power method, as well as the matches returned by semidefinite relaxation [65, 66]. The computation runtime is reported in Table 1. The numerical results demonstrate that the projected power method is significantly faster than SDP, while achieving a joint matching performance as competitive as SDP.

## 5 Preliminaries and notation

Starting from this section, we turn to the analyses of the main results. Before proceeding, we gather a few preliminary facts and notations that will be useful throughout.

### 5.1 Projection onto the standard simplex

Firstly, our algorithm involves projection onto the standard simplex  $\Delta$ . In light of this, we single out several elementary facts concerning  $\Delta$  and  $\mathcal{P}_\Delta$  as follows.

**Fact 1.** Suppose that  $\mathbf{a} = [a_1, \dots, a_m]^\top$  obeys  $\mathbf{a} + \mathbf{e}_l \in \Delta$  for some  $1 \leq l \leq m$ . Then  $\|\mathbf{a}\| \leq \sqrt{2}$ .

*Proof.* The feasibility condition requires  $0 \leq \sum_{i:i \neq l} a_i = -a_l \leq 1$  and  $a_i \geq 0$  for all  $i \neq l$ . Therefore, it is easy to check that  $\sum_{i:i \neq l} a_i^2 \leq \sum_{i:i \neq l} a_i \leq 1$ , and hence  $\|\mathbf{a}\|^2 = a_l^2 + \sum_{i:i \neq l} a_i^2 \leq 2$ .  $\square$

**Fact 2.** For any vector  $\mathbf{v} \in \mathbb{R}^m$  and any value  $\delta$ , one has

$$\mathcal{P}_\Delta(\mathbf{v}) = \mathcal{P}_\Delta(\mathbf{v} + \delta \mathbf{1}). \quad (46)$$

*Proof.* For any  $\mathbf{x} \in \Delta$ ,

$$\|\mathbf{v} + \delta \mathbf{1} - \mathbf{x}\|^2 = \|\mathbf{v} - \mathbf{x}\|^2 + \delta^2 \|\mathbf{1}\|^2 + 2\delta(\mathbf{v} - \mathbf{x})^\top \mathbf{1} = \|\mathbf{v} - \mathbf{x}\|^2 + \delta^2 n + 2\delta(\mathbf{v}^\top \mathbf{1} - 1).$$

<sup>14</sup>[https://github.com/huangqx/CSP\\_Codes](https://github.com/huangqx/CSP_Codes)

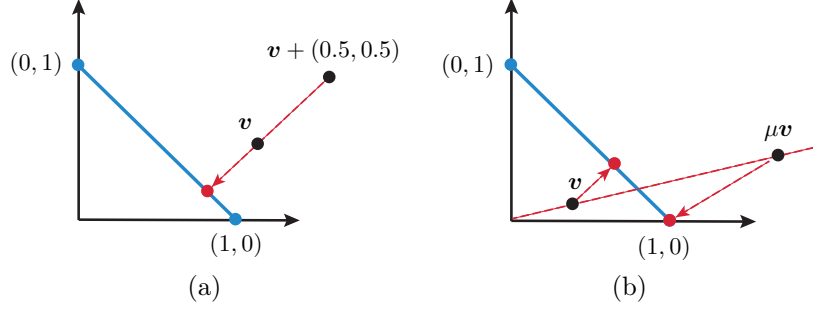


Figure 7: Illustration of Facts 2-3 on a two-dimensional standard simplex. (a) For any  $\mathbf{v}$ ,  $\mathcal{P}_\Delta(\mathbf{v}) = \mathcal{P}_\Delta(\mathbf{v} + 0.5 \cdot \mathbf{1})$ ; (b) for an vector  $\mathbf{v} \in \mathbb{R}^2$  obeying  $v_1 > v_2$ , one has  $\mathcal{P}_\Delta(\mu \mathbf{v}) = \mathbf{e}_1$  when  $\mu$  is sufficiently large.

Hence,  $\mathcal{P}_\Delta(\mathbf{v} + \delta \mathbf{1}) = \arg \min_{\mathbf{x} \in \Delta} \|\mathbf{v} + \delta \mathbf{1} - \mathbf{x}\|^2 = \arg \min_{\mathbf{x} \in \Delta} \|\mathbf{v} - \mathbf{x}\|^2 = \mathcal{P}_\Delta(\mathbf{v})$ .  $\square$

**Fact 3.** For any non-zero vector  $\mathbf{v} = [v_i]_{1 \leq i \leq m}$ , let  $v_{(1)}$  and  $v_{(2)}$  be its largest and second largest entries, respectively. Suppose  $v_j = v_{(1)}$ . If  $\mu > 1/(v_{(1)} - v_{(2)})$ , then

$$\mathcal{P}_\Delta(\mu \mathbf{v}) = \mathbf{e}_j. \quad (47)$$

*Proof.* By convexity of  $\Delta$ , we have  $\mathcal{P}_\Delta(\mu \mathbf{v}) = \mathbf{e}_j$  if and only if, for any  $\mathbf{x} = [x_i]_{1 \leq i \leq m} \in \Delta$ ,

$$(\mathbf{x} - \mathbf{e}_j)^\top (\mu \mathbf{v} - \mathbf{e}_j) \leq 0.$$

Since  $\mathbf{x}^\top \mathbf{v} = x_j v_{(1)} + \sum_{i:i \neq j} x_i v_i \leq x_j v_{(1)} + v_{(2)} \sum_{i:i \neq j} x_i = x_j v_{(1)} + v_{(2)}(1 - x_j)$ , we see that

$$(\mathbf{x} - \mathbf{e}_j)^\top (\mu \mathbf{v} - \mathbf{e}_j) \leq (1 - x_j)(1 - \mu(v_{(1)} - v_{(2)})) \leq 0.$$

$\square$

In words, Fact 2 claims that a global offset does not alter the projection  $\mathcal{P}_\Delta(\cdot)$ , while Fact 3 reveals that a large scaling factor  $\mu$  results in sufficient separation between the largest entry and the remaining ones. See Fig. 7 for a graphical illustration.

## 5.2 Properties of the likelihood ratios

Next, we study the log-likelihood ratio statistics. The first result makes a connection between the KL divergence and other properties of the log-likelihood ratio. Here and below, for any two distributions  $P$  and  $Q$  supported on  $\mathcal{Y}$ , the total variation distance between them is defined by  $\text{TV}(P, Q) := \frac{1}{2} \sum_{y \in \mathcal{Y}} |P(y) - Q(y)|$ .

**Lemma 2.** (1) Consider two probability distributions  $P$  and  $Q$  on a finite set  $\mathcal{Y}$ . Then

$$\left| \log \frac{Q(y)}{P(y)} \right| \leq \frac{2\text{TV}(P, Q)}{\min\{P(y), Q(y)\}} \leq \frac{\sqrt{2\text{KL}(P \parallel Q)}}{\min\{P(y), Q(y)\}}. \quad (48)$$

(2) In addition, if both  $\max_{y \in \mathcal{Y}} \frac{P(y)}{Q(y)} \leq \kappa_0$  and  $\max_{y \in \mathcal{Y}} \frac{Q(y)}{P(y)} \leq \kappa_0$  hold, then

$$\mathbb{E}_{y \sim P} \left[ \left( \log \frac{P(y)}{Q(y)} \right)^2 \right] \leq 2\kappa_0^2 \text{KL}(Q \parallel P). \quad (49)$$

*Proof.* See Appendix C.  $\square$

In particular, when  $\text{KL}(P \parallel Q)$  is small, one almost attains equality in (49), as stated below.

**Lemma 3.** Consider two probability distributions  $P$  and  $Q$  on a finite set  $\mathcal{Y}$  such that  $P(y)$  and  $Q(y)$  are both bounded away from zero. If  $\text{KL}(P \parallel Q) \leq \varepsilon$ , then one has

$$\text{KL}(P \parallel Q) = \frac{1 + \zeta_1(\varepsilon)}{2} \mathbf{Var}_{y \sim P} \left[ \log \frac{P(y)}{Q(y)} \right] \quad (50)$$

$$\text{and} \quad \text{H}^2(P, Q) = \frac{1 + \zeta_2(\varepsilon)}{4} \text{KL}(P \parallel Q), \quad (51)$$

where  $\zeta_1$  and  $\zeta_2$  are functions satisfying  $|\zeta_1(\varepsilon)|, |\zeta_2(\varepsilon)| \leq c_0 \sqrt{\varepsilon}$  for some universal constant  $c_0 > 0$ .

*Proof.* See Appendix D.  $\square$

### 5.3 Block random matrices

Additionally, the data matrix  $\mathbf{L}$  is assumed to have independent blocks. It is thus crucial to control the fluctuation of such i.i.d. random block matrices, for which the following lemma proves useful.

**Lemma 4.** Let  $\mathbf{M} := [\mathbf{M}_{i,j}]_{1 \leq i,j \leq n}$  be any random symmetric block matrix, where  $\{\mathbf{M}_{i,j} \in \mathbb{R}^{m \times m} \mid i \geq j\}$  are independently generated. Suppose that  $m = O(\text{poly}(n))$ ,  $\mathbb{E}[\mathbf{M}_{i,j}] = \mathbf{0}$ ,  $\max_{i,j} \|\mathbf{M}_{i,j}\| \leq K$ , and  $\mathbb{P}\{\mathbf{M}_{i,j} = \mathbf{0}\} = p_{\text{obs}}$  for some  $p_{\text{obs}} \gtrsim \log n/n$ . Then with probability exceeding  $1 - O(n^{-10})$ ,

$$\|\mathbf{M}\| \lesssim K \sqrt{np_{\text{obs}}}. \quad (52)$$

*Proof.* See Appendix E.  $\square$

Lemma 4 immediately leads to an upper estimate on the fluctuations of  $\mathbf{L}$  and  $\mathbf{L}^{\text{debias}}$ .

**Lemma 5.** Suppose  $m = O(\text{poly}(n))$ , and define  $\left\| \log \frac{P_0}{P_l} \right\|_1 := \sum_y \left| \log \frac{P_0(y)}{P_l(y)} \right|$ . If  $p_{\text{obs}} \gtrsim \log n/n$ , then with probability exceeding  $1 - O(n^{-10})$ , the matrices  $\mathbf{L}$  and  $\mathbf{L}^{\text{debias}}$  given respectively in (7) and (12) satisfy

$$\|\mathbf{L} - \mathbb{E}[\mathbf{L}]\| = \|\mathbf{L}^{\text{debias}} - \mathbb{E}[\mathbf{L}^{\text{debias}}]\| \lesssim \left( \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1 \right) \sqrt{np_{\text{obs}}}. \quad (53)$$

*Proof.* See Appendix F.  $\square$

### 5.4 Other notation

Several other notations are in order. For any vector  $\mathbf{h} = [\mathbf{h}_i]_{1 \leq i \leq n} \in \mathbb{R}^{mn}$  with  $\mathbf{h}_i \in \mathbb{R}^m$ , we denote by  $h_{i,j}$  the  $j$ th component of  $\mathbf{h}_i$ . For any  $m \times n$  matrix  $\mathbf{A} = [a_{i,j}]_{1 \leq i \leq m, 1 \leq j \leq n}$  and any matrix  $\mathbf{B}$ , the Kronecker product  $\mathbf{A} \otimes \mathbf{B}$  is defined as

$$\mathbf{A} \otimes \mathbf{B} := \begin{bmatrix} a_{1,1}\mathbf{B} & \cdots & a_{1,n}\mathbf{B} \\ \vdots & \ddots & \vdots \\ a_{m,1}\mathbf{B} & \cdots & a_{m,n}\mathbf{B} \end{bmatrix}.$$

Throughout this paper,  $\log(\cdot)$  represents the natural logarithm.

## 6 Iterative stage

We establish the performance guarantees of our two-stage algorithm in a reverse order. Specifically, we demonstrate in this section that the iterative refinement stage achieves exact recovery, provided that the initial guess is reasonably close to the truth. The analysis for the initialization is deferred to Section 7.

## 6.1 Error contraction

This section mainly consists in establishing the following claim, which concerns error contraction of iterative refinement in the presence of an appropriate initial guess.

**Theorem 8.** *Under the conditions of Theorem 3 or Theorem 7, there exist some absolute constants  $0 < \rho, c_1 < 1$  such that with probability exceeding  $1 - O(n^{-5})$ ,*

$$\|\mathcal{P}_{\Delta^n}(\mu \mathbf{L} \mathbf{z}) - \mathbf{x}\|_{*,0} \leq \rho \min \left\{ \|\mathbf{z} - \mathbf{x}\|_{*,0}, \|\mathbf{z} - \mathbf{x}\|^2 \right\} \quad (54)$$

holds simultaneously for all  $\mathbf{z} \in \Delta^n$  obeying<sup>15</sup>

$$\min \left\{ \frac{\|\mathbf{z} - \mathbf{x}\|}{\|\mathbf{x}\|}, \frac{\|\mathbf{z} - \mathbf{x}\|_{*,0}}{\|\mathbf{x}\|_{*,0}} \right\} \leq 0.49 \frac{\text{KL}_{\min}}{\text{KL}_{\max}}, \quad (55)$$

provided that

$$\mu > c_5 / (np_{\text{obs}} \text{KL}_{\min}) \quad (56)$$

for some sufficiently large constant  $c_5 > 0$ .

At each iteration, the PPM produces a more accurate estimate as long as the iterates  $\{\mathbf{z}^{(t)}\}$  stay within a reasonable neighborhood surrounding  $\mathbf{x}$ . Here and below, we term this neighborhood a *basin of attraction*. In fact, if the initial guess  $\mathbf{z}^{(0)}$  successfully lands within this basin, then the subsequent iterates will never jump out of it. To see this, observe that for any  $\mathbf{z}^{(t)}$  obeying

$$\frac{\|\mathbf{z}^{(t)} - \mathbf{x}\|}{\|\mathbf{x}\|} \leq 0.49 \frac{\text{KL}_{\min}}{\text{KL}_{\max}} \quad \text{or} \quad \frac{\|\mathbf{z}^{(t)} - \mathbf{x}\|_{*,0}}{\|\mathbf{x}\|_{*,0}} \leq 0.49 \frac{\text{KL}_{\min}}{\text{KL}_{\max}},$$

the inequality (54) implies error contraction

$$\|\mathbf{z}^{(t+1)} - \mathbf{x}\|_{*,0} \leq \rho \|\mathbf{z}^{(t)} - \mathbf{x}\|_{*,0}.$$

Moreover, since  $\|\mathbf{x}\|_{*,0} = \|\mathbf{x}\|^2 = n$ , one has

$$\frac{\|\mathbf{z}^{(t+1)} - \mathbf{x}\|_{*,0}}{\|\mathbf{x}\|_{*,0}} < \min \left\{ \frac{\|\mathbf{z}^{(t)} - \mathbf{x}\|_{*,0}}{\|\mathbf{x}\|_{*,0}}, \frac{\|\mathbf{z}^{(t)} - \mathbf{x}\|^2}{\|\mathbf{x}\|^2} \right\} \leq \min \left\{ \frac{\|\mathbf{z}^{(t)} - \mathbf{x}\|_{*,0}}{\|\mathbf{x}\|_{*,0}}, \frac{\|\mathbf{z}^{(t)} - \mathbf{x}\|}{\|\mathbf{x}\|} \right\} \leq 0.49 \frac{\text{KL}_{\min}}{\text{KL}_{\max}},$$

precluding the possibility that  $\mathbf{z}^{(t+1)}$  leaves the basin. As a result, invoking the preceding theorem iteratively we arrive at

$$\|\mathbf{z}^{(t)} - \mathbf{x}\|_{*,0} \leq \rho^t \|\mathbf{z}^{(0)} - \mathbf{x}\|_{*,0},$$

indicating that the estimation error reduces to zero within at most logarithmic iterations.

**Remark 5.** In fact, the contraction rate  $\rho$  can be as small as  $O\left(\frac{1}{np_{\text{obs}} \text{KL}_{\min}}\right) = O\left(\frac{1}{\log n}\right)$  in the scenario considered in Theorem 3 or  $O\left(\max_l \left\| \log \frac{P_0}{P_l} \right\|_1^2 / (np_{\text{obs}} \text{KL}_{\min}^2)\right)$  in the case studied in Theorem 7.

Furthermore, we emphasize that Theorem 8 is a uniform result, namely, it holds simultaneously for all  $\mathbf{z}$  within the basin, regardless of whether  $\mathbf{z}$  is independent of the data  $\{y_{i,j}\}$  or not. Consequently, the theory and the analyses remain valid for other initialization schemes that can produce a suitable first guess.

The rest of the section is thus devoted to establishing Theorem 8. The proofs for the two scenarios—the fixed  $m$  case and the large  $m$  case—follow almost identical arguments, and hence we shall combine the analyses together.

<sup>15</sup>The numerical constant 0.49 is arbitrary and can be replaced by any other constant in  $(0, 0.5)$ .

## 6.2 Analysis

We outline the key steps for the proof of Theorem 8. Before continuing, it is helpful to introduce additional assumptions and notation that will be used throughout. From now on, we will assume  $\mathbf{x}_i = \mathbf{e}_1$  for all  $1 \leq i \leq n$  without loss of generality. We shall denote  $\mathbf{h} = \{\mathbf{h}_i\}_{1 \leq i \leq n} \in \mathbb{R}^{nm}$  and  $\mathbf{w} = \{\mathbf{w}_i\}_{1 \leq i \leq n} \in \mathbb{R}^{nm}$  as

$$\mathbf{h} := \mathbf{z} - \mathbf{x} \quad \text{and} \quad \mathbf{w} := \mathbf{L}\mathbf{z}, \quad (57)$$

and set

$$k := \|\mathbf{h}\|_{*,0}; \quad (58)$$

$$\epsilon := \|\mathbf{h}\|/\|\mathbf{x}\| = \|\mathbf{h}\|/\sqrt{n}; \quad (59)$$

$$k^* := \min \left\{ \|\mathbf{h}\|_{*,0}, \frac{\|\mathbf{h}\|^2}{\|\mathbf{x}\|^2} \|\mathbf{x}\|_{*,0} \right\} = \min \{k, \epsilon^2 n\}. \quad (60)$$

One of the key metrics that will play an important role in our proof is the following separation measure

$$\mathcal{S}(\mathbf{a}) := \min_{2 \leq l \leq m} (a_1 - a_l) \quad (61)$$

defined for any vector  $\mathbf{a} = [a_l]_{1 \leq l \leq m} \in \mathbb{R}^m$ . This metric is important because, by Fact 3, the projection of a block  $\mu \mathbf{w}_i$  onto the standard simplex  $\Delta$  returns the correct solution—that is,  $\mathcal{P}_\Delta(\mu \mathbf{w}_i) = \mathbf{e}_1$ —as long as  $\mathcal{S}(\mathbf{w}_i) > 0$  and  $\mu$  is sufficiently large. As such, our aim is to show that the vector  $\mathbf{w}$  given in (57) obeys

$$\mathcal{S}(\mathbf{w}_i) > 0.01 n p_{\text{obs}} \text{KL}_{\min} \quad \forall i \in \mathcal{I} \subseteq \{1, \dots, n\}, \quad (62)$$

for some index set  $\mathcal{I}$  of size  $|\mathcal{I}| \geq n - \rho k^*$ , where  $0 < \rho < 1$  is bounded away from 1 (which will be specified later). This taken collectively with Fact 3 implies  $\mathcal{P}_\Delta(\mu \mathbf{w}_i) = \mathbf{x}_i = \mathbf{e}_1$  for every  $i \in \mathcal{I}$  and, as a result,

$$\|\mathcal{P}_{\Delta^n}(\mu \mathbf{w}) - \mathbf{x}\|_{*,0} \leq \sum_{i \notin \mathcal{I}} \|\mathcal{P}_\Delta(\mu \mathbf{w}_i) - \mathbf{x}_i\|_0 = n - |\mathcal{I}| \leq \rho k^* = \rho \min \left\{ \|\mathbf{h}\|_{*,0}, \frac{\|\mathbf{h}\|}{\|\mathbf{x}\|} \|\mathbf{x}\|_{*,0} \right\},$$

provided that the scaling factor obeys  $\mu > 100/(n p_{\text{obs}} \text{KL}_{\min})$ .

We will organize the proof of the claim (62) based on the size / block sparsity of  $\mathbf{h}$ , leaving us with two separate regimes to deal with:

- The *large-error regime* in which

$$\xi < \min \left\{ \frac{k}{n}, \epsilon \right\} \leq 0.49 \frac{\text{KL}_{\min}}{\text{KL}_{\max}}; \quad (63)$$

- The *small-error regime* in which

$$\min \left\{ \frac{k}{n}, \epsilon \right\} \leq \xi. \quad (64)$$

Here, one can take  $\xi > 0$  to be any small positive constant.

(1) **Large-error regime.** Suppose that  $\mathbf{z}$  falls within the regime (63). In order to control  $\mathcal{S}(\mathbf{w}_i)$ , we decompose  $\mathbf{w} = \mathbf{L}\mathbf{z}$  into a few terms that are easier to work with. Specifically, setting  $\bar{\mathbf{h}} := \frac{1}{n} \sum_{i=1}^n \mathbf{h}_i$  and  $\tilde{\mathbf{L}} := \mathbf{L} - \mathbb{E}[\mathbf{L}]$  we can expand

$$\mathbf{w} = \mathbf{L}\mathbf{z} = \left( \mathbb{E}[\mathbf{L}] + \tilde{\mathbf{L}} \right) (\mathbf{x} + \mathbf{h}) = \underbrace{\mathbb{E}[\mathbf{L}]\mathbf{x}}_{:=\mathbf{t}} + \underbrace{\mathbb{E}[\mathbf{L}]\mathbf{h} + \tilde{\mathbf{L}}\mathbf{x} + \tilde{\mathbf{L}}\mathbf{h}}_{:=\mathbf{r}}. \quad (65)$$

This allows us to lower bound the separation for the  $i$ th component by

$$\mathcal{S}(\mathbf{w}_i) = \mathcal{S}(\mathbf{t}_i + \mathbf{r}_i) \geq \mathcal{S}(\mathbf{t}_i) + \mathcal{S}(\mathbf{r}_i) = \mathcal{S}(\mathbf{t}_i) + \min_{2 \leq l \leq m} (r_{i,1} - r_{i,l}) \geq \mathcal{S}(\mathbf{t}_i) - 2\|\mathbf{r}_i\|_\infty. \quad (66)$$

With this in mind, attention naturally turns to controlling  $\mathcal{S}(\mathbf{t}_i)$  and  $\|\mathbf{r}_i\|_\infty$ .

The first quantity  $\mathcal{S}(\mathbf{t}_i)$  admits a closed-form expression. From (35) and (65) one sees that

$$\mathbf{t}_i = p_{\text{obs}} \left( \sum_{j:j \neq i} \mathbf{K} \mathbf{x}_j \right) = p_{\text{obs}} (n-1) \mathbf{K} \mathbf{e}_1.$$

It is self-evident that  $\mathcal{S}(\mathbf{K} \mathbf{e}_1) = \text{KL}_{\min}$ , giving the formula

$$\mathcal{S}(\mathbf{t}_i) = p_{\text{obs}} (n-1) \mathcal{S}(\mathbf{K} \mathbf{e}_1) = p_{\text{obs}} (n-1) \text{KL}_{\min}. \quad (67)$$

We are now faced with the problem of estimating  $\|\mathbf{r}_i\|_\infty$ . To this end, we make the following observation, which holds uniformly over all  $\mathbf{z}$  residing within this regime:

**Lemma 6.** *Consider the regime (63). Suppose  $m = O(\text{poly}(n))$ ,  $p_{\text{obs}} \gtrsim \log n/n$ , and*

$$\frac{\text{KL}_{\max}^2}{\left\{ \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1 \right\}^2} > \frac{c_2}{np_{\text{obs}}} \quad (68)$$

for some sufficiently large constant  $c_2 > 0$ . With probability exceeding  $1 - O(n^{-10})$ , the index set

$$\mathcal{I} := \left\{ 1 \leq i \leq n \mid \|\mathbf{r}_i\|_\infty \leq np_{\text{obs}} \text{KL}_{\max} \min \left\{ \frac{k}{n}, \epsilon \right\} + \alpha np_{\text{obs}} \text{KL}_{\max} \right\} \quad (69)$$

has cardinality exceeding  $n - \rho k^*$  for some  $0 < \rho < 1$  bounded away from 1, where  $k^* = \min \{k, \epsilon^2 n\}$  and  $\alpha > 0$  is some arbitrarily small constant.

In particular, if  $m$  is fixed and if Assumption 1 holds, then (68) can be replaced by

$$\text{KL}_{\max} > \frac{c_4}{np_{\text{obs}}} \quad (70)$$

for some sufficiently large constant  $c_4 > 0$ .

*Proof.* See Appendix G. □

Combining Lemma 6 with the preceding bounds (66) and (67), we obtain

$$\mathcal{S}(\mathbf{w}_i) \geq (n-1) p_{\text{obs}} \text{KL}_{\min} - 2np_{\text{obs}} \text{KL}_{\max} \min \left\{ \frac{k}{n}, \epsilon \right\} - 2\alpha np_{\text{obs}} \text{KL}_{\max} \quad (71)$$

$$> 0.01 np_{\text{obs}} \text{KL}_{\min} \quad (72)$$

for all  $i \in \mathcal{I}$  as given in (69), provided that (i)  $\min \left\{ \frac{k}{n}, \epsilon \right\} \leq 0.49 \frac{\text{KL}_{\min}}{\text{KL}_{\max}}$ , (ii)  $\text{KL}_{\max}/\text{KL}_{\min}$  is bounded, (iii)  $\alpha$  is sufficiently small, and (iv)  $n$  is sufficiently large. This concludes the treatment for the large-error regime.

(2) **Small-error regime.** We now turn to the second regime obeying (64). Similarly, we find it convenient to decompose  $\mathbf{w}$  as

$$\mathbf{w} = \mathbf{L} \mathbf{x} + \mathbf{L} \mathbf{h} = \underbrace{\mathbf{L} \mathbf{x}}_{:=\mathbf{s}} + \underbrace{\mathbb{E}[\mathbf{L}] \mathbf{h} + \tilde{\mathbf{L}} \mathbf{h}}_{:=\mathbf{q}}, \quad (73)$$

and then lower bound the separation measure by controlling  $\mathbf{s}_i$  and  $\mathbf{q}_i$  separately, i.e.

$$\mathcal{S}(\mathbf{w}_i) \geq \mathcal{S}(\mathbf{s}_i) + \mathcal{S}(\mathbf{q}_i) \geq \mathcal{S}(\mathbf{s}_i) - 2 \|\mathbf{q}_i\|_\infty. \quad (74)$$

We start by obtaining a uniform control on the separation of all components of  $\mathbf{s}$ :

**Lemma 7.** *Suppose that Assumption 1 holds and that  $p_{\text{obs}} > c_0 \log n/n$  for some sufficiently large constant  $c_0 > 0$ .*

(1) *Fix  $m > 0$ , and let  $\zeta > 0$  be any sufficiently small constant. Under Condition (28), one has*

$$\mathcal{S}(\mathbf{s}_i) > \zeta np_{\text{obs}} \text{KL}_{\min}, \quad 1 \leq i \leq n \quad (75)$$

with probability exceeding  $1 - C_6 \exp\{-c_6 \zeta \log(nm)\} - c_7 n^{-10}$ , where  $C_6, c_6, c_7 > 0$  are some absolute constants.

(2) There exist some constants  $c_4, c_5, c_6 > 0$  such that

$$\mathcal{S}(\mathbf{s}_i) > c_4 n p_{\text{obs}} \text{KL}_{\min}, \quad 1 \leq i \leq n \quad (76)$$

with probability  $1 - O(m^{-10} n^{-10})$ , provided that

$$\frac{\text{KL}_{\min}^2}{\max_{0 \leq l < m} \text{Var}_{y \sim P_0} \left[ \log \frac{P_0(y)}{P_l(y)} \right]} \geq \frac{c_5 \log(mn)}{n p_{\text{obs}}} \quad \text{and} \quad \text{KL}_{\min} \geq \frac{c_6 \left\{ \max_{l,y} \left| \log \frac{P_0(y)}{P_l(y)} \right| \right\} \log(mn)}{n p_{\text{obs}}}. \quad (77)$$

*Proof.* See Appendix H.  $\square$

The next step comes down to controlling  $\|\mathbf{q}_i\|_\infty$ . This can be accomplished using similar argument as for Lemma 6, as summarized below.

**Lemma 8.** *Consider the regime (64). Then Lemma 6 continues to hold if  $\mathbf{r}$  is replaced by  $\mathbf{q}$ .*

Putting Assumption 2, the inequality (74), and Lemma 8 together yields

$$\mathcal{S}(\mathbf{w}_i) \geq \mathcal{S}(\mathbf{s}_i) - 2n p_{\text{obs}} \text{KL}_{\max} \min \left\{ \frac{k}{n}, \epsilon \right\} - 2\alpha n p_{\text{obs}} \text{KL}_{\max} \geq \mathcal{S}(\mathbf{s}_i) - 4\tilde{\xi} n p_{\text{obs}} \text{KL}_{\min} \quad (78)$$

for all  $i \in \mathcal{I}$  with high probability, provided that  $\min \left\{ \frac{k}{n}, \epsilon \right\} \leq \tilde{\xi}$  and  $\alpha \leq \tilde{\xi}$  for some constant  $\tilde{\xi} > 0$ . Picking  $\tilde{\xi}$  to be sufficiently small and applying Lemma 7, we arrive at (62).

To summarize, we have established the claim (62)—and hence the error contraction—as long as (a)  $m$  is fixed and Condition (28) is satisfied, or (b) the conditions (68) and (77) hold. Interestingly, one can simplify Case (b) when  $p_{\text{obs}} \gtrsim \log^5 n/n$ , leading to a matching condition to Theorem 7.

**Lemma 9.** *Suppose  $m \gtrsim \log n$ ,  $m = \text{poly}(n)$ , and  $p_{\text{obs}} \geq c_6 \log^5 n/n$  for some sufficiently large constant  $c_6 > 0$ . The inequalities (68) and (77) hold under Condition (44) in addition to Assumptions 1-2.*

*Proof.* See Appendix I.  $\square$

### 6.3 Choice of the scaling factor $\mu$

So far we have proved the result under the scaling factor condition (56) given in Theorem 8. To conclude the analysis for Theorem 3 and Theorem 7, it remains to convert it to conditions in terms of the singular value  $\sigma_i(\cdot)$ .

To begin with, it follows from (34) that

$$\mathbf{L} = \mathbb{E}[\mathbf{L}] + (\mathbf{L} - \mathbb{E}[\mathbf{L}]) = p_{\text{obs}} \mathbf{1}\mathbf{1}^\top \otimes \mathbf{K} - p_{\text{obs}} \mathbf{I}_n \otimes \mathbf{K} + (\mathbf{L} - \mathbb{E}[\mathbf{L}]),$$

leading to an upper estimate

$$\begin{aligned} \sigma_i(\mathbf{L}) &\leq \sigma_i(p_{\text{obs}} \mathbf{1}\mathbf{1}^\top \otimes \mathbf{K}) + \|p_{\text{obs}} \mathbf{I}_n \otimes \mathbf{K}\| + \|\mathbf{L} - \mathbb{E}[\mathbf{L}]\| \\ &= n p_{\text{obs}} \sigma_i(\mathbf{K}) + p_{\text{obs}} \|\mathbf{K}\| + \|\mathbf{L} - \mathbb{E}[\mathbf{L}]\|. \end{aligned} \quad (79)$$

Since  $\mathbf{K}$  is circulant, its eigenvalues are given by

$$\lambda_l = \sum_{i=0}^{m-1} (-\text{KL}_i - \mathcal{H}(P_0)) \exp\left(j \frac{2\pi i l}{m}\right), \quad 0 \leq l < m$$

with  $j = \sqrt{-1}$ . In fact, except for  $\lambda_0$ , one can simplify

$$\lambda_l = - \sum_{i=0}^{m-1} \text{KL}_i \exp\left(j \frac{2\pi i l}{m}\right), \quad 1 \leq l < m,$$

which are eigenvalues of  $\mathbf{K}^0$  as well. This enables the upper bounds

$$\sigma_2(\mathbf{K}) \leq \sqrt{\sum_{j=1}^{m-1} \lambda_j^2} \leq \|\mathbf{K}^0\|_F \leq m\mathbf{KL}_{\max} \lesssim m\mathbf{KL}_{\min}, \quad (80)$$

$$\sigma_m(\mathbf{K}) \leq \sqrt{\frac{1}{m-1} \sum_{j=1}^{m-1} \lambda_j^2} \leq \sqrt{\frac{1}{m-1} \|\mathbf{K}^0\|_F^2} \leq \sqrt{\frac{m^2 \mathbf{KL}_{\max}^2}{m-1}} \leq \sqrt{2m} \mathbf{KL}_{\max} \lesssim \sqrt{m} \mathbf{KL}_{\min}, \quad (81)$$

where the last inequalities of both (80) and (81) follow from Assumption 2. In addition, it is immediate to see that (79) and (81) remain valid if we replace  $\mathbf{L}$  with  $\mathbf{L}^{\text{debias}}$  and take  $\mathbf{K} = \frac{1}{p_{\text{obs}}} \mathbb{E}[\mathbf{L}_{i,j}^{\text{debias}}]$  ( $i \neq j$ ) instead.

To bound the remaining terms on the right-hand side of (79), we divide into two separate cases:

(i) When  $m$  is fixed, it follows from (35) that

$$\|\mathbf{K}\| \leq \|\mathbf{K}_0\| + m\mathcal{H}(P_0) \leq m\mathbf{KL}_{\max} + m \log m = O(1).$$

When combined with Lemma 5, this yields

$$p_{\text{obs}}\|\mathbf{K}\| + \|\tilde{\mathbf{L}}\| \lesssim p_{\text{obs}} + \left( \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1 \right) \sqrt{np_{\text{obs}}} \lesssim p_{\text{obs}} + \sqrt{np_{\text{obs}} \mathbf{KL}_{\max}}, \quad (82)$$

where the last inequality follows from (147). Putting this together with (79) and (80) and using Assumption 2, we get

$$\sigma_2(\mathbf{L}) \lesssim np_{\text{obs}} m \mathbf{KL}_{\min} + p_{\text{obs}} + \sqrt{np_{\text{obs}} \mathbf{KL}_{\max}} \asymp np_{\text{obs}} \mathbf{KL}_{\min}. \quad (83)$$

Thus, one would satisfy (56) by taking  $\mu \geq c_{12}/\sigma_2(\mathbf{L})$  for some sufficiently large  $c_{12} > 0$ .

(ii) When  $m = O(\text{poly}(n))$ , we consider  $\mathbf{L}^{\text{debias}}$  and set  $\mathbf{K} = \frac{1}{p_{\text{obs}}} \mathbb{E}[\mathbf{L}_{i,j}^{\text{debias}}]$ . According to (134), one has  $\|\mathbf{L}_{i,j}^{\text{debias}}\| \leq \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1$ , thus indicating that

$$p_{\text{obs}}\|\mathbf{K}\| = \|\mathbb{E}[\mathbf{L}_{i,j}^{\text{debias}}]\| \leq \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1. \quad (84)$$

This together with Lemma 5 gives

$$p_{\text{obs}}\|\mathbf{K}\| + \|\mathbf{L}^{\text{debias}} - \mathbb{E}[\mathbf{L}^{\text{debias}}]\| \lesssim \left( \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1 \right) (1 + \sqrt{np_{\text{obs}}}) \quad (85)$$

$$\asymp \left( \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1 \right) \sqrt{np_{\text{obs}}} \lesssim np_{\text{obs}} \mathbf{KL}_{\min} \quad (86)$$

under the condition (44). Putting all of this together, we derive

$$\sigma_m(\mathbf{L}^{\text{debias}}) \lesssim \sqrt{mnp_{\text{obs}} \mathbf{KL}_{\min}} + np_{\text{obs}} \mathbf{KL}_{\min} \asymp \sqrt{mnp_{\text{obs}} \mathbf{KL}_{\min}},$$

thus justifying (56) as long as  $\mu \geq c_{12}\sqrt{m}/\sigma_m(\mathbf{L}^{\text{debias}})$  for some sufficiently large constant  $c_{12} > 0$ .

## 6.4 Consequences for random corruption models

Having obtained a qualitative behavior of the iterative stage for general models, we can now specialize it to the random corruption model (3). Before continuing, it is straightforward to compute two metrics:

$$\mathbf{KL}_{\min} = \mathbf{KL}_{\max} = \left( \pi_0 + \frac{1-\pi_0}{m} \right) \log \frac{\pi_0 + \frac{1-\pi_0}{m}}{\frac{1-\pi_0}{m}} + \frac{1-\pi_0}{m} \log \frac{\frac{1-\pi_0}{m}}{\pi_0 + \frac{1-\pi_0}{m}} = \pi_0 \log \frac{1 + (m-1)\pi_0}{1-\pi_0};$$

$$\text{and} \quad \left\| \log \frac{P_0}{P_l} \right\|_1 = 2 \left| \log \frac{\pi_0 + \frac{1-\pi_0}{m}}{\frac{1-\pi_0}{m}} \right| = 2 \log \frac{1 + (m-1)\pi_0}{1-\pi_0}.$$



1. When  $m$  is fixed and  $\pi_0$  is small, it is not hard to see that

$$\text{KL}_{\min} = \text{KL}_{\max} \approx m\pi_0^2,$$

which taken collectively with (28) leads to (20).

2. When  $m \gtrsim \log n$  and  $m = O(\text{poly}(n))$ , the condition (44) reduces to  $\pi_0^2 \gtrsim 1/(np_{\text{obs}})$ , which coincides with (41). In fact, one can also easily verify (77) under Condition (41), assuming that  $p_{\text{obs}} \gtrsim \log^2 n/n$ . This improves slightly upon the condition  $p_{\text{obs}} \gtrsim \log^5 n/n$  required in the general theorem.

Next, we demonstrate that the algorithm with the input matrix (11) undergoes the same trajectory as the version using (7). To avoid confusion, we shall let  $\mathbf{L}^{\text{rcm}}$  denote the matrix (11), and set  $\mathbf{w}^{\text{rcm}} := \mathbf{L}^{\text{rcm}} \mathbf{z}$ . As discussed before, there are some constants  $a > 0$  and  $b$  such that  $\mathbf{L}_{i,j}^{\text{rcm}} = a\mathbf{L}_{i,j} + b\mathbf{1}\mathbf{1}^\top$  for all  $(i, j) \in \Omega$ , indicating that

$$\mathbf{w}_i^{\text{rcm}} = a\mathbf{w}_i + \tilde{b}_i \mathbf{1}, \quad 1 \leq i \leq n$$

for some numerical values  $\{\tilde{b}_i\}$ . In view of Fact 2, the projection  $\mathcal{P}_\Delta$  remains unchanged up to global shift. This justifies the equivalence between the two input matrices when running Algorithm 1.

Finally, one would have to adjust the scaling factor accordingly. It is straightforward to show that the scaling factor condition (56) can be translated into  $\mu > c_5/(np_{\text{obs}}\pi_0)$  when the input matrix (11) is employed. Observe that (79) continues to hold as long as we set  $\mathbf{K} = \pi_0 \mathbf{I}_m + (1 - \pi_0) \mathbf{1}\mathbf{1}^\top$ . We can also verify that

$$\sigma_2(\mathbf{K}) = \pi_0, \quad \|\mathbf{K}\| \leq 1, \quad \text{and} \quad \|\mathbf{L} - \mathbb{E}[\mathbf{L}]\| \lesssim \sqrt{np_{\text{obs}}}, \quad (87)$$

where the last inequality follows from Lemma 4. These taken collectively with (79) lead to

$$\sigma_2(\mathbf{L}) \lesssim np_{\text{obs}}\pi_0 + p_{\text{obs}}\pi_0 + \sqrt{np_{\text{obs}}} \lesssim np_{\text{obs}}\pi_0 \quad (88)$$

under the condition (41). This justifies the choice  $\mu \gtrsim 1/\sigma_2(\mathbf{L})$  as desired.

## 7 Spectral initialization

We move on to the performance of spectral initialization by establishing the theorem below. Similar to the definition (18), we introduce the counterpart of  $\ell_2$  distance modulo the global offset as

$$\text{dist}(\hat{\mathbf{x}}, \mathbf{x}) := \min_{0 \leq l < m} \|\hat{\mathbf{x}} - \text{shift}_l(\mathbf{x})\|.$$

**Theorem 9.** Fix  $\delta > 0$ , and suppose that  $p_{\text{obs}} \gtrsim \log n/n$ . Under Assumptions 1-2, there are some universal constants  $c_1, c_2, c_3 > 0$  such that with probability at least  $1 - \xi$ , the initial estimate  $\mathbf{z}^{(0)}$  of Algorithms 1 obeys

$$\text{dist}(\mathbf{z}^{(0)}, \mathbf{x}) \leq \delta \|\mathbf{x}\| \quad \text{and} \quad \text{MCR}(\mathbf{z}^{(0)}, \mathbf{x}) \leq \delta^2/2 \quad (89)$$

in the following scenarios:

- (i) the random corruption model with  $m = O(\text{poly}(n))$ , provided that  $\mathbf{L}$  is given by (11) and that

$$\pi_0 \geq \frac{c_1}{\delta \sqrt{\xi}} \frac{1}{\sqrt{p_{\text{obs}} n}} \quad (90)$$

- (iii) the general model with fixed  $m$ , provided that  $\mathbf{L}$  is given by (7) and that

$$\text{KL}_{\min} \geq \frac{c_2}{\delta^2 \xi} \frac{1}{p_{\text{obs}} n}. \quad (91)$$

- (iii) the general model with  $m = O(\text{poly}(n))$ , provided that  $\mathbf{L}$  is replaced by  $\mathbf{L}^{\text{debias}}$  (given in (12)) and that

$$\frac{\text{KL}_{\min}^2}{\max_{1 \leq l < m} \left\| \log \frac{P_0}{P_l} \right\|_1^2} \geq \frac{c_3}{\delta^2 \xi} \frac{1}{p_{\text{obs}} n}. \quad (92)$$

The main reason for spectral initialization to succeed is that the low-rank approximation of  $\mathbf{L}$  (resp.  $\mathbf{L}^{\text{debias}}$ ) produce a decent estimate of  $\mathbb{E}[\mathbf{L}]$  (resp.  $\mathbb{E}[\mathbf{L}^{\text{debias}}]$ ) and, as discussed before,  $\mathbb{E}[\mathbf{L}]$  (resp.  $\mathbb{E}[\mathbf{L}^{\text{debias}}]$ ) reveals the structure of the truth. In what follows, we will first prove the result for general  $\mathbf{L}$ , and then specialize it to the three choices considered in the theorem. As usual, we suppose without loss of generality that  $\mathbf{x}_i = \mathbf{e}_1$ ,  $1 \leq i \leq n$ .

To begin with, we set  $\mathbf{K} = \frac{1}{p_{\text{obs}}} \mathbb{E}[\mathbf{L}_{i,j}]$  ( $i \neq j$ ) as before and write

$$\mathbf{L} = p_{\text{obs}} \mathbf{1}\mathbf{1}^\top \otimes \mathbf{K} - p_{\text{obs}} \mathbf{I}_n \otimes \mathbf{K} + (\mathbf{L} - \mathbb{E}[\mathbf{L}]), \quad (93)$$

where the first term on the right-hand side of (93) has rank at most  $m$ . If we let  $\hat{\mathbf{L}}$  be the best rank- $m$  approximation of  $\mathbf{L}$ , then matrix perturbation theory gives

$$\|\hat{\mathbf{L}} - p_{\text{obs}} \mathbf{1}\mathbf{1}^\top \otimes \mathbf{K}\| \leq \|p_{\text{obs}} \mathbf{I}_n \otimes \mathbf{K}\| + \|\mathbf{L} - \mathbb{E}[\mathbf{L}]\| = p_{\text{obs}} \|\mathbf{K}\| + \|\mathbf{L} - \mathbb{E}[\mathbf{L}]\| := \gamma_{\mathbf{L}}. \quad (94)$$

This together with the facts  $\text{rank}(\hat{\mathbf{L}}) \leq m$  and  $\text{rank}(\mathbf{K}) \leq m$  gives

$$\|\hat{\mathbf{L}} - p_{\text{obs}} \mathbf{1}\mathbf{1}^\top \otimes \mathbf{K}\|_{\text{F}}^2 \leq 2m \|\hat{\mathbf{L}} - p_{\text{obs}} \mathbf{1}\mathbf{1}^\top \otimes \mathbf{K}\|^2 \leq 2m \gamma_{\mathbf{L}}^2.$$

Further, let  $\mathbf{K}_{:,1}$  (resp.  $\mathbf{1} \otimes \mathbf{K}_{:,1}$ ) be the first column of  $\mathbf{K}$  (resp.  $\mathbf{1} \otimes \mathbf{K}$ ). When  $\hat{\mathbf{z}}$  is taken to be a random column of  $\hat{\mathbf{L}}$ , it is straightforward to verify that

$$\mathbb{E} [\text{dist}^2(\hat{\mathbf{z}}, p_{\text{obs}} \mathbf{1} \otimes \mathbf{K}_{:,1})] \leq \frac{1}{nm} \mathbb{E} [\|\hat{\mathbf{L}} - p_{\text{obs}} \mathbf{1}\mathbf{1}^\top \otimes \mathbf{K}\|_{\text{F}}^2] \leq \frac{2\gamma_{\mathbf{L}}^2}{n}.$$

Apply Markov's inequality to deduce that, with probability at least  $1 - \xi$ ,

$$\text{dist}(\hat{\mathbf{z}}, p_{\text{obs}} \mathbf{1} \otimes \mathbf{K}_{:,1}) \leq \frac{\sqrt{2}\gamma_{\mathbf{L}}}{\sqrt{\xi n}}. \quad (95)$$

For simplicity of presentation, we shall assume

$$\|\hat{\mathbf{z}} - p_{\text{obs}} \mathbf{1} \otimes \mathbf{K}_{:,1}\| = \text{dist}(\hat{\mathbf{z}}, p_{\text{obs}} \mathbf{1} \otimes \mathbf{K}_{:,1})$$

from now on. We shall pay particular attention to the index set

$$\mathcal{J} := \left\{ i \in [n] \mid \|\hat{\mathbf{z}}_i - p_{\text{obs}} \mathbf{K}_{:,1}\| \leq \frac{\sqrt{2}}{\delta \sqrt{n}} \|\hat{\mathbf{z}} - p_{\text{obs}} \mathbf{1} \otimes \mathbf{K}_{:,1}\| = \frac{\sqrt{2}}{\delta \sqrt{n}} \text{dist}(\hat{\mathbf{z}}, p_{\text{obs}} \mathbf{1} \otimes \mathbf{K}_{:,1}) \right\},$$

which clearly satisfies  $|\mathcal{J}| \geq (1 - \delta^2/2)n$  due to the fact that

$$\|\hat{\mathbf{z}} - p_{\text{obs}} \mathbf{1} \otimes \mathbf{K}_{:,1}\|^2 \geq \sum_{i \notin \mathcal{J}} \|\hat{\mathbf{z}}_i - p_{\text{obs}} \mathbf{K}_{:,1}\|^2 > (n - |\mathcal{J}|) \cdot \frac{2}{\delta^2 n} \|\hat{\mathbf{z}} - p_{\text{obs}} \mathbf{1} \otimes \mathbf{K}_{:,1}\|^2.$$

The  $\ell_\infty$  error in each block can also be bounded by

$$\|\hat{\mathbf{z}}_i - p_{\text{obs}} \mathbf{K}_{:,1}\|_\infty \leq \|\hat{\mathbf{z}}_i - p_{\text{obs}} \mathbf{K}_{:,1}\| \leq \frac{\sqrt{2}}{\delta \sqrt{n}} \text{dist}(\hat{\mathbf{z}}, p_{\text{obs}} \mathbf{1} \otimes \mathbf{K}_{:,1}), \quad i \in \mathcal{J}. \quad (96)$$

If the above  $\ell_\infty$  error is sufficiently small for each  $i \in \mathcal{J}$ , then the projection operation recovers the truth for all blocks falling in  $\mathcal{J}$ . Specifically, adopting the separation measure  $\mathcal{S}(\cdot)$  defined in (61), we obtain

$$\mathcal{S}(\hat{\mathbf{z}}_i) \geq \mathcal{S}(p_{\text{obs}} \mathbf{K}_{:,1}) + \mathcal{S}(\hat{\mathbf{z}}_i - p_{\text{obs}} \mathbf{K}_{:,1}) \geq p_{\text{obs}} \mathcal{S}(\mathbf{K}_{:,1}) - 2\|\hat{\mathbf{z}}_i - p_{\text{obs}} \mathbf{K}_{:,1}\|_\infty.$$

If  $\|\hat{\mathbf{z}}_i - p_{\text{obs}} \mathbf{K}_{:,1}\|_\infty \leq c_5 p_{\text{obs}} \mathcal{S}(\mathbf{K}_{:,1})$ ,  $i \in \mathcal{J}$  for some constant  $0 < c_5 < 1/2$ , then it follows from Fact 3 that

$$\mathbf{z}_i^{(0)} = \mathcal{P}_\Delta(\mu_0 \hat{\mathbf{z}}_i) = \mathbf{e}_1, \quad i \in \mathcal{J}$$

as long as  $\mu_0 > \frac{1}{(1-2c_5)p_{\text{obs}}\mathcal{S}(\mathbf{K}_{:,1})}$ . This taken collectively with Fact 1 further reveals that

$$\begin{cases} \text{dist}(\mathbf{z}^{(0)}, \mathbf{x}) \leq \sqrt{2} \cdot \sqrt{n - |\mathcal{J}|} \leq \delta\sqrt{n} = \delta\|\mathbf{x}\| \\ \text{MCR}(\mathbf{z}^{(0)}, \mathbf{x}) \leq (n - |\mathcal{J}|)/n \leq \delta^2/2 \end{cases} \quad (97)$$

as claimed. As a result, everything boils down to proving that

$$\|\hat{\mathbf{z}}_i - p_{\text{obs}}\mathbf{K}_{:,1}\|_\infty \leq c_5 p_{\text{obs}} \mathcal{S}(\mathbf{K}_{:,1}), \quad i \in \mathcal{J}. \quad (98)$$

In view of (95) and (96), this condition (98) would hold if

$$\gamma_{\mathbf{L}} \leq c_6 \delta \sqrt{\xi n p_{\text{obs}} \mathcal{S}(\mathbf{K}_{:,1})} \quad (99)$$

for some sufficiently small constant  $c_6 > 0$ .

To finish up, we establish (99) for the three scenarios considered Theorem 9.

- (i) *The random corruption model with  $\mathbf{L}$  given by (11).* Simple calculation gives  $\mathcal{S}(\mathbf{K}_{:,1}) = \pi_0$ , and it follows from (87) that  $\gamma_{\mathbf{L}} \lesssim \sqrt{np_{\text{obs}}}$ . Thus, the condition (99) would hold under the assumption (90).
- (ii) *The general model with fixed  $m$  and with  $\mathbf{L}$  given by (7).* Observe that  $\mathcal{S}(\mathbf{K}_{:,1}) = \text{KL}_{\min}$ , and recall from (82) that  $\gamma_{\mathbf{L}} \lesssim p_{\text{obs}} + \sqrt{np_{\text{obs}}\text{KL}_{\max}}$ . Thus, we necessarily have (99) under the condition (91) and Assumption 2.
- (iii) *The general model with  $m = O(\text{poly}(n))$  and with  $\mathbf{L}$  replaced by  $\mathbf{L}^{\text{debias}}$ , in which  $\mathcal{S}(\mathbf{K}_{:,1}) = \text{KL}_{\min}$ .* It has been shown in (86) that

$$\gamma_{\mathbf{L}^{\text{debias}}} \lesssim \left( \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1 \right) \sqrt{np_{\text{obs}}}.$$

As a consequence, we establish (99) under the assumption (92).

Finally, repeating the analyses for the scaling factor in Section 6 finishes the proof.

## 8 Minimax lower bound

This section proves the minimax lower bound in Theorem 4. Once this is done, we can apply it to the random corruption model, which immediately establishes Theorem 2 using exactly the same calculation as in Section 6.4.

To prove Theorem 4, it suffices to analyze the ML rule, which minimizes the Bayesian error probability when we impose a uniform prior over all possible inputs. Before continuing, we provide an asymptotic estimate on the tail exponent of the likelihood ratio test, which proves crucial in bounding the probability of error of ML decoding.

**Lemma 10.** *Let  $\{P_n\}_{n \geq 1}$  and  $\{Q_n\}_{n \geq 1}$  be two sequences of probability measures on a fixed finite set  $\mathcal{Y}$ , where  $\min_{n,y} P_n(y)$  and  $\min_{n,y} Q_n(y)$  are both bounded away from 0. Let  $\{y_{j,n} : 1 \leq j \leq n\}_{n \geq 1}$  be a triangular array of independent random variables such that  $y_{j,n} \sim P_n$ . If we define*

$$\mu_n := \text{KL}(P_n \parallel Q_n) \quad \text{and} \quad \sigma_n^2 := \text{Var}_{y \sim P_n} \left[ \log \frac{Q_n(y)}{P_n(y)} \right],$$

then for any given constant  $\tau > 0$ ,

$$\mathbb{P} \left\{ \frac{1}{n} \sum_{j=1}^n \log \frac{Q_n(y_{j,n})}{P_n(y_{j,n})} + \mu_n > \tau \mu_n \right\} = \exp \left\{ - (1 + o_n(1)) n \frac{\tau^2 \mu_n^2}{2\sigma_n^2} \right\} = \exp \left\{ - \frac{1 + o_n(1)}{4} \tau^2 n \mu_n \right\} \quad (100)$$

and

$$\mathbb{P} \left\{ \frac{1}{n} \sum_{j=1}^n \log \frac{Q_n(y_{j,n})}{P_n(y_{j,n})} + \mu_n < -\tau \mu_n \right\} = \exp \left\{ - (1 + o_n(1)) n \frac{\tau^2 \mu_n^2}{2\sigma_n^2} \right\} = \exp \left\{ - \frac{1 + o_n(1)}{4} \tau^2 n \mu_n \right\} \quad (101)$$

hold as long as  $\frac{\mu_n^2}{\sigma_n^2} \asymp \frac{\log n}{n}$  and  $\mu_n \asymp \frac{\log n}{n}$ .

*Proof.* This lemma is a consequence of the moderate deviation theory. See Appendix J.  $\square$

**Remark 6.** The asymptotic limits presented in Lemma 10 correspond to the Gaussian tail, meaning that some sort of the central limit theorem holds in this regime.

We shall now freeze the input to be  $x_1 = \dots = x_n = 1$  and consider the conditional error probability of the ML rule. Without loss of generality, we assume  $P_{0,n}$  and  $P_{1,n}$  are minimally separated, namely,

$$\text{KL}(P_{0,n} \parallel P_{1,n}) = \text{KL}_{\min,n}. \quad (102)$$

In what follows, we will suppress the dependence on  $n$  whenever clear from the context, and let  $\mathbf{x}^{\text{ML}}$  represent the ML estimate. We claim that it suffices to prove Theorem 4 for the boundary regime where

$$\text{KL}_{\min} \asymp \frac{\log n}{np_{\text{obs}}} \quad \text{and} \quad \text{KL}_{\min} \leq \frac{3.99 \log n}{np_{\text{obs}}}. \quad (103)$$

In fact, suppose instead that the error probability

$$P_e := \mathbb{P}\{\text{MCR}(\mathbf{x}^{\text{ML}}, \mathbf{x}) > 0\}$$

tends to one in the regime (103) but is bounded away from one when  $\text{KL}_{\min} = o\left(\frac{\log n}{np_{\text{obs}}}\right)$ . Then this indicates that, in the regime (103), one can always add extra noise<sup>16</sup> to  $y_{i,j}$  to decrease  $\text{KL}_{\min}$  while significantly improving the success probability, which results in contradiction. Moreover, when  $P_n(y)$  and  $Q_n(y)$  are both bounded away from 0, it follows from Lemma 3 that

$$\text{Var}_{y \sim P_{0,n}} \left[ \log \frac{P_{1,n}(y)}{P_{0,n}(y)} \right] \asymp \text{KL}_{\min} \asymp \frac{\log n}{np_{\text{obs}}}, \quad \text{and} \quad \frac{\text{KL}_{\min}^2}{\text{Var}_{y \sim P_{0,n}} \left[ \log \frac{P_{1,n}(y)}{P_{0,n}(y)} \right]} \asymp \frac{\log n}{np_{\text{obs}}}. \quad (104)$$

Consider a set  $\mathcal{V}_1 = \{1, \dots, \delta n\}$  for some small constant  $\delta > 0$ . We first single out a subset  $\mathcal{V}_2 \subseteq \mathcal{V}_1$  such that the local likelihood ratio score—when restricted to samples over the subgraph induced by  $\mathcal{V}_1$ —is sufficiently large. More precisely, we take

$$\mathcal{V}_2 := \left\{ i \in \mathcal{V}_1 \mid \sum_{j: j \in \mathcal{V}_1} \log \frac{P_1(y_{i,j})}{P_0(y_{i,j})} > -2\delta np_{\text{obs}} \text{KL}(P_0 \parallel P_1) \right\};$$

here and throughout, we set  $\frac{P_1(y_{i,j})}{P_0(y_{i,j})} = 1$  for any  $(i, j) \notin \Omega$  for notational simplicity. Recall that for each  $i$ , the ML rule favors  $P_1$  (resp.  $x_i = 2$ ) against  $P_0$  (resp.  $x_i = 1$ ) if and only if

$$\sum_j \log \frac{P_1(y_{i,j})}{P_0(y_{i,j})} > 0,$$

which would happen if

$$\sum_{j: j \in \mathcal{V}_1} \log \frac{P_1(y_{i,j})}{P_0(y_{i,j})} > -2\delta np_{\text{obs}} \text{KL}(P_0 \parallel P_1) \quad \text{and} \quad \sum_{j: j \notin \mathcal{V}_1} \log \frac{P_1(y_{i,j})}{P_0(y_{i,j})} > 2\delta np_{\text{obs}} \text{KL}(P_0 \parallel P_1).$$

Thus, conditional on  $\mathcal{V}_2$  we can lower bound the probability of error by

$$\begin{aligned} P_e &\geq \mathbb{P} \left\{ \exists i \in \mathcal{V}_2 : \sum_j \log \frac{P_1(y_{i,j})}{P_0(y_{i,j})} > 0 \right\} \\ &\geq \mathbb{P} \left\{ \exists i \in \mathcal{V}_2 : \sum_{j: j \in \mathcal{V}_1} \log \frac{P_1(y_{i,j})}{P_0(y_{i,j})} > -2\delta np_{\text{obs}} \text{KL}(P_0 \parallel P_1) \quad \text{and} \quad \sum_{j: j \notin \mathcal{V}_1} \log \frac{P_1(y_{i,j})}{P_0(y_{i,j})} > 2\delta np_{\text{obs}} \text{KL}(P_0 \parallel P_1) \right\} \\ &= \mathbb{P} \left\{ \exists i \in \mathcal{V}_2 : \sum_{j: j \notin \mathcal{V}_1} \log \frac{P_1(y_{i,j})}{P_0(y_{i,j})} > 2\delta np_{\text{obs}} \text{KL}(P_0 \parallel P_1) \right\}, \end{aligned} \quad (105)$$

---

<sup>16</sup>For instance, we can let  $\tilde{y}_{i,j} = \begin{cases} y_{i,j}, & \text{with probability } \rho_0 \\ \text{Unif}(m), & \text{else} \end{cases}$  with  $\rho_0$  controlling the noise level.

where the last identity comes from the definition of  $\mathcal{V}_2$ .

We pause to remark on why (105) facilitates analysis. To begin with,  $\mathcal{V}_2$  depends only on those samples lying within the subgraph induced by  $\mathcal{V}_1$ , and is thus independent of  $\sum_{j \notin \mathcal{V}_1} \log \frac{P_1(y_{i,j})}{P_0(y_{i,j})}$ . More importantly, the scores  $\left\{s_i := \sum_{j: j \notin \mathcal{V}_1} \log \frac{P_1(y_{i,j})}{P_0(y_{i,j})}\right\}$  are statistically independent across all  $i \in \mathcal{V}_1$  as they rely on distinct samples. These allow to derive that, conditional on  $\mathcal{V}_2$ ,

$$(105) = 1 - \prod_{i \in \mathcal{V}_2} \left(1 - \mathbb{P} \left\{ \sum_{j: j \notin \mathcal{V}_1} \log \frac{P_1(y_{i,j})}{P_0(y_{i,j})} > 2\delta n p_{\text{obs}} \text{KL}(P_0 \parallel P_1) \right\}\right) \\ \geq 1 - \left(1 - \exp \left\{ - (1 + o(1)) (1 + 2\delta)^2 n p_{\text{obs}} \frac{\text{KL}_{\min}}{4} \right\}\right)^{|\mathcal{V}_2|} \quad (106)$$

$$\geq 1 - \exp \left\{ - |\mathcal{V}_2| \exp \left\{ - (1 + o(1)) (1 + 2\delta)^2 n p_{\text{obs}} \frac{\text{KL}_{\min}}{4} \right\} \right\}, \quad (107)$$

where the last line results from the elementary inequality  $1 - x \leq e^{-x}$ . To see why (106) holds, we note that according to the Chernoff bound, the number of samples linking each  $i$  and  $\bar{\mathcal{V}}_1$  (i.e.  $|\{j \notin \mathcal{V}_1 : (i, j) \in \Omega\}|$ ) is at most  $n p_{\text{obs}}$  with high probability, provided that (i)  $\frac{n p_{\text{obs}}}{\log n}$  is sufficiently large, and (ii)  $n$  is sufficiently large. These taken collectively with Lemma 10 and (104) yield

$$\mathbb{P} \left\{ \sum_{j: j \notin \mathcal{V}_1} \log \frac{P_1(y_{i,j})}{P_0(y_{i,j})} > 2\delta n p_{\text{obs}} \text{KL}(P_0 \parallel P_1) \right\} \geq \exp \left\{ - (1 + o(1)) (1 + 2\delta)^2 n p_{\text{obs}} \frac{\text{KL}_{\min}}{4} \right\},$$

thus justifying (106).

To establish Theorem 4, we would need to show that (107) (and hence  $P_e$ ) is lower bounded by  $1 - o(1)$  or, equivalently,

$$|\mathcal{V}_2| \exp \left\{ - (1 + o(1)) (1 + 2\delta)^2 n p_{\text{obs}} \text{KL}_{\min} / 4 \right\} \rightarrow \infty.$$

This condition would hold if

$$(1 + 2\delta)^2 n p_{\text{obs}} \text{KL}_{\min} / 4 < (1 - \delta) \log n \quad (108)$$

and

$$|\mathcal{V}_2| = (1 - o(1)) |\mathcal{V}_1| = (1 - o(1)) \delta n, \quad (109)$$

since

$$\begin{aligned} |\mathcal{V}_2| \exp \left\{ - (1 + 2\delta)^2 n p_{\text{obs}} \text{KL}_{\min} / 4 \right\} &\geq |\mathcal{V}_2| \exp \left\{ - (1 - \delta) \log n \right\} \\ &\geq \delta \exp \left\{ \log n - (1 - \delta) \log n \right\} \\ &\geq \delta n^\delta \rightarrow \infty. \end{aligned}$$

The first condition (108) is a consequence from (37) as long as  $\delta$  is sufficiently small. It remains to verify the second condition (109).

When  $n p_{\text{obs}} > c_0 \log n$  for some sufficiently large constant  $c_0 > 0$ , each  $i \in \mathcal{V}_1$  is connected to at least  $(1 - \delta) p_{\text{obs}} |\mathcal{V}_1|$  vertices in  $\mathcal{V}_1$  with high probability, meaning that the number of random variables involved in the sum  $\sum_{j: j \in \mathcal{V}_1} \log \frac{P_1(y_{i,j})}{P_0(y_{i,j})}$  concentrates around  $|\mathcal{V}_1| p_{\text{obs}}$ . Lemma 10 thus implies that

$$\begin{aligned} \mathbb{P} \left\{ \sum_{j: j \in \mathcal{V}_1} \log \frac{P_1(y_{i,j})}{P_0(y_{i,j})} \leq -2\delta n p_{\text{obs}} \text{KL}(\mathbb{P}_0 \parallel \mathbb{P}_1) \right\} &= \mathbb{P} \left\{ \sum_{j: j \in \mathcal{V}_1} \log \frac{P_1(y_{i,j})}{P_0(y_{i,j})} \leq -2|\mathcal{V}_1| p_{\text{obs}} \text{KL}(\mathbb{P}_0 \parallel \mathbb{P}_1) \right\} \\ &\leq \exp \left\{ -c_4 (|\mathcal{V}_1| p_{\text{obs}}) \text{KL}_{\min} \right\} \leq \exp \left\{ -c_5 (\delta n p_{\text{obs}}) \text{KL}_{\min} \right\} \end{aligned}$$

for some constants  $c_4, c_5 > 0$ , provided that  $n$  is sufficiently large. This gives rise to an upper bound

$$\begin{aligned} \mathbb{E} [|\mathcal{V}_1 \setminus \mathcal{V}_2|] &= |\mathcal{V}_1| \cdot \mathbb{P} \left\{ \sum_{j: j \in \mathcal{V}_1} \log \frac{P_1(y_{i,j})}{P_0(y_{i,j})} \leq -2\delta n \text{KL}(P_0 \parallel P_1) \right\} \\ &\leq \delta n \cdot \exp \left\{ -c_5 \delta n p_{\text{obs}} \text{KL}_{\min} \right\} \\ &\stackrel{(i)}{=} \delta n \cdot n^{-\Theta(\delta)} = o(n), \end{aligned}$$

where (i) arises from the condition (103). As a result, Markov's inequality implies that with probability approaching one,

$$|\mathcal{V}_1 \setminus \mathcal{V}_2| = o(n)$$

or, equivalently,  $|\mathcal{V}_2| = (1 - o(1)) \delta n$ . This finishes the proof of Theorem 4.

## 9 Discussion

We have developed an efficient nonconvex paradigm to compute the MLE for a class of discrete assignment problems. There are numerous questions we leave open that might be interesting for future investigation. For instance, it can be seen from Fig. 1 and Fig. 2 that the algorithm returns reasonably good estimates even when we are below the information limits. A natural question is this: how can we characterize the accuracy of the algorithm if one is satisfied with approximate solutions? In addition, this work assumes the index set  $\Omega$  of the pairwise samples are drawn uniformly at random. Depending on the application scenarios, we might encounter other measurement patterns that cannot be modeled in this random manner. Can we determine the performance of the algorithm for more general sampling set  $\Omega$ ? Moreover, the log-likelihood functions we incorporate in the data matrix  $\mathbf{L}$  might be imperfect. Further study could help understand the stability of the algorithm in the presence of model mismatch.

Returning to Assumption 2, we remark that this assumption is imposed primarily out of computational concern. In fact,  $\text{KL}_{\max}/\text{KL}_{\min}$  being exceedingly large might actually be a favorable case from an information theoretic viewpoint, as it indicates that the hypothesis corresponding to  $\text{KL}_{\max}$  is much easier to preclude compared to other hypotheses. It would be interesting to establish rigorously the performance of the PPM without this assumption and, in case it becomes suboptimal, how shall we modify the algorithm so as to be more adaptive to the most general class of noise models.

Moving beyond joint alignment, we are interested in seeing the potential benefits of the PPM on other discrete problems. For instance, the joint alignment problem falls under the category of maximum *a posteriori* (MAP) inference in a discrete Markov random field, which spans numerous applications including segmentation, object detection, error correcting codes, and so on [74–76]. Specifically, consider  $n$  discrete variables  $x_i \in [m]$ ,  $1 \leq i \leq n$ . We are given a set of unitary potential functions (or prior distributions)  $\{\psi_i(x_i)\}_{1 \leq i \leq n}$  as well as a collection of pairwise potential functions (or likelihood functions)  $\{\psi_{i,j}(x_i, x_j)\}_{(i,j) \in \mathcal{G}}$  over some graph  $\mathcal{G}$ . The goal is to compute the MAP assignment

$$\mathbf{x}_{\text{MAP}} := \arg \max_{\mathbf{z}} \prod_{i=1}^n \psi_i(z_i) \prod_{(i,j) \in \mathcal{G}} \psi_{i,j}(z_i, z_j) = \arg \max_{\mathbf{z}} \left\{ \sum_{i=1}^n \log \psi_i(z_i) + \sum_{(i,j) \in \mathcal{G}} \log \psi_{i,j}(z_i, z_j) \right\}. \quad (110)$$

Similar to (6) and (7), one can introduce the vector  $\mathbf{z}_i = \mathbf{e}_j$  to represent  $z_i = j$ , and use a matrix  $\mathbf{L}_{i,j} \in \mathbb{R}^{m \times m}$  to encode each pairwise log-potential function  $\log \psi_{i,j}(\cdot, \cdot)$ ,  $(i, j) \in \mathcal{G}$ . The unitary potential function  $\psi_i(\cdot)$  can also be encoded by a diagonal matrix  $\mathbf{L}_{i,i} \in \mathbb{R}^{m \times m}$

$$(\mathbf{L}_{i,i})_{\alpha,\alpha} = \log \psi_i(\alpha), \quad 1 \leq \alpha \leq m \quad (111)$$

so that  $\log \psi(z_i) = \mathbf{z}_i^\top \mathbf{L}_{i,i} \mathbf{z}_i$ . This enables a quadratic form representation of MAP estimation:

$$\begin{aligned} & \text{maximize}_{\mathbf{z}} && \mathbf{z}^\top \mathbf{L} \mathbf{z} \\ & \text{subject to} && \mathbf{z}_i \in \{\mathbf{e}_1, \dots, \mathbf{e}_m\}, \quad 1 \leq i \leq n. \end{aligned} \quad (112)$$

As such, we expect the PPM to be effective in solving many instances of such MAP inference problems. One of the key questions amounts to finding an appropriate initialization that allows efficient exploitation of the unitary prior belief  $\psi_i(\cdot)$ . We leave this for future work.

## Acknowledgements

E. C. is partially supported by the Math + X Award from the Simons Foundation. Y. C. is supported by the same award. We thank Qixing Huang for motivating discussions about the join image alignment problem. Y. Chen is grateful to Qixing Huang, Leonidas Guibas, and Nan Hu for helpful discussions about joint graph matching.

## A Proof of Theorem 5

We will concentrate on proving the case where Assumption 1 is violated. Set  $y_{\max} := \arg \max_y P_0(y)$  and  $y_{\min} := \arg \min_y P_0(y)$ . Since  $m$  is fixed, it is seen that  $P_0(y_{\max}) \asymp 1$ . We also have  $P_0(y_{\min}) \rightarrow 0$ . Denoting by  $\nu := \frac{P_0(y_{\max})}{P_0(y_{\min})}$  the dynamic range of  $P_0$  and introducing the metric

$$d(P \parallel Q) := \sum_y P(y) \left| \log \frac{P(y)}{Q(y)} \right|,$$

we obtain

$$\max_l d(P_0 \parallel P_l) = \max_l \sum_y P_0(y) \left| \log \frac{P_0(y)}{P_l(y)} \right| \asymp P_0(y_{\max}) \left| \log \frac{P_0(y_{\max})}{P_0(y_{\min})} \right| \asymp \log \nu \gg 1.$$

The elementary inequality  $\text{KL}(P \parallel Q) \leq d(P \parallel Q)$ , together with the second Pinsker's inequality  $\text{KL}(P \parallel Q) + \sqrt{2\text{KL}(P \parallel Q)} \geq d(P \parallel Q)$  [77, Lemma 2.5], reveals that

$$\text{KL}_{\max} \asymp \max_l d(P_0 \parallel P_l) \asymp \log \nu.$$

Making use of Assumption 2 we get

$$d(P_0 \parallel P_l) \asymp \text{KL}_{\min} \asymp \text{KL}_{\max} \asymp \log \nu, \quad 1 \leq l < m. \quad (113)$$

Next, for each  $1 \leq l < m$  we single out an element  $y_l := \arg \max_y P_0(y) \left| \log \frac{P_0(y)}{P_l(y)} \right|$ . This element is important because, when  $m$  is fixed,

$$d(P_0 \parallel P_l) \asymp P_0(y_l) \left| \log \frac{P_0(y_l)}{P_l(y_l)} \right| \lesssim P_0(y_l) \log \nu.$$

As a result, (113) would only happen if  $P_0(y_l) \asymp 1$  and  $\frac{P_0(y_l)}{P_l(y_l)} \rightarrow \infty$ .

We are now ready to prove the theorem. With  $P_l$  replaced by  $\tilde{P}_l$  as defined in (40), one has

$$\text{KL}(\tilde{P}_0 \parallel \tilde{P}_l) \geq 2\text{TV}^2(\tilde{P}_0, \tilde{P}_l) \geq \frac{1}{2}(\tilde{P}_0(y_l) - \tilde{P}_l(y_l))^2 \asymp (P_0(y_l) - P_l(y_l))^2 \asymp P_0^2(y_l) \asymp 1,$$

which exceeds the threshold  $4.01 \log n / (np_{\text{obs}})$  as long as  $p_{\text{obs}} \geq c_1 \log n / n$  for some sufficiently large constant  $c_1$ . In addition, since  $\tilde{P}_0(y)$  is bounded away from both 0 and 1, it is easy to see that  $\text{KL}(\tilde{P}_0 \parallel \tilde{P}_l) \lesssim 1$  for any  $l$  and, hence, Assumption 2 remains valid. Invoking Theorem 3 concludes the proof.

## B Proof of Lemma 1

(1) It suffices to prove the case where  $\text{KL}_{\min} = \text{KL}(P_0 \parallel P_1)$ . Suppose that

$$|P_0(y_0) - P_l(y_0)| = \max_{j,y} |P_0(y) - P_j(y)|$$

for some  $0 \leq l, y_0 < m$ . In view of Pinsker's inequality [77, Lemma 2.5],

$$\begin{aligned} \text{KL}(P_0 \parallel P_1) &\geq 2\text{TV}^2(P_0, P_1) = 2 \left( \frac{1}{2} \sum_{y=0}^{m-1} |P_0(y) - P_1(y)| \right)^2 = \frac{1}{2} \left( \sum_{y=0}^{m-1} |P_0(y) - P_0(y-1)| \right)^2 \\ &\geq \frac{1}{2} |P_0(y_0) - P_0(y_0-l)|^2 = \frac{1}{2} \max_{j,y} |P_0(y) - P_j(y)|^2. \end{aligned} \quad (114)$$

On the other hand, for any  $1 \leq j < m$ ,

$$\begin{aligned} \text{KL}(P_0 \parallel P_j) &\stackrel{(a)}{\leq} \chi^2(P_0 \parallel P_j) := \sum_y \frac{(P_0(y) - P_j(y))^2}{P_j(y)} \stackrel{(b)}{\asymp} \sum_y (P_0(y) - P_j(y))^2 \\ &\stackrel{(c)}{\lesssim} \max_{j,y} |P_0(y) - P_j(y)|^2, \end{aligned} \quad (115)$$

where (a) comes from [78, Eqn. (5)], (b) is a consequence from Assumption 1, and (c) follows since  $m$  is fixed. Combining (114) and (115) establishes  $\text{KL}_{\min} \asymp \text{KL}_{\max}$ .

(2) Suppose that  $P_0(y^*) = \min_y P_0(y)$  for  $y^* = \lfloor m/2 \rfloor$ , and that  $\text{KL}(P_0 \parallel P_l) = \text{KL}_{\min}$  for some  $0 < l \leq y^*$ . Applying Pinsker's inequality again gives

$$\begin{aligned} \text{KL}(P_0 \parallel P_l) &\geq 2\text{TV}^2(P_0, P_l) = \frac{1}{2} \left( \sum_{y=0}^{m-1} |P_l(y) - P_0(y)| \right)^2 = \frac{1}{2} \left( \sum_{y=0}^{m-1} |P_0(y-l) - P_0(y)| \right)^2 \\ &\geq \frac{1}{2} \left( |P_0(y^*) - P_0(y^* - l)| + \sum_{k=1}^{\lceil y^*/l \rceil - 1} |P_0(kl) - P_0((k-1)l)| \right)^2 \\ &\geq \frac{1}{2} \left( |P_0(y^*) - P_0((\lceil y^*/l \rceil - 1)l)| + \sum_{k=1}^{\lceil y^*/l \rceil - 1} |P_0(kl) - P_0((k-1)l)| \right)^2 \end{aligned} \quad (116)$$

$$\geq \frac{1}{2} |P_0(y^*) - P_0(0)|^2 = \frac{1}{2} \max_y |P_0(0) - P_0(y)|^2 = \frac{1}{2} \max_{j,y} |P_0(y) - P_j(y)|^2, \quad (117)$$

where (116) follows from the unimodality assumption, and the last line results from the facts  $P_0(0) = \max_{j,y} P_j(y)$  and  $P_0(y^*) = \min_{j,y} P_j(y)$ . These taken collectively with (115) finish the proof.

## C Proof of Lemma 2

(1) Since  $\log(1+x) \leq x$  for any  $x \geq 0$ , we get

$$\left| \log \frac{Q(y)}{P(y)} \right| = \begin{cases} \log \left( 1 + \frac{Q(y) - P(y)}{P(y)} \right), & \text{if } Q(y) \geq P(y) \\ \log \left( 1 + \frac{P(y) - Q(y)}{Q(y)} \right), & \text{else} \end{cases} \leq \frac{|Q(y) - P(y)|}{\min\{P(y), Q(y)\}} \quad (118)$$

$$\leq \frac{2\text{TV}(P, Q)}{\min\{P(y), Q(y)\}} \leq \frac{\sqrt{2\text{KL}(P \parallel Q)}}{\min\{P(y), Q(y)\}}, \quad (119)$$

where the last inequality comes from Pinsker's inequality.

(2) Using the inequality (118) once again as well as the definition of  $\kappa_0$ , we obtain

$$\begin{aligned} \mathbb{E}_{y \sim P} \left[ \left( \log \frac{P(y)}{Q(y)} \right)^2 \right] &= \sum_y P(y) \left( \log \frac{P(y)}{Q(y)} \right)^2 \\ &\leq \sum_y P(y) \frac{|P(y) - Q(y)|^2}{\min\{P^2(y), Q^2(y)\}} = \sum_y \frac{|P(y) - Q(y)|^2}{\min\left\{ \frac{P(y)}{Q(y)}, \frac{Q(y)}{P(y)} \right\} Q(y)} \end{aligned} \quad (120)$$

$$\leq \kappa_0 \sum_y \frac{(Q(y) - P(y))^2}{Q(y)}. \quad (121)$$

Note that  $\sum \frac{(Q(y) - P(y))^2}{Q(y)}$  is exactly the  $\chi_2$  divergence between  $P$  and  $Q$ , which satisfies [79, Proposition 2]

$$\chi_2(P \parallel Q) \leq 2\kappa_0 \text{KL}(P \parallel Q). \quad (122)$$

Substitution into (121) concludes the proof.

## D Proof of Lemma 3

For notational simplicity, let

$$\mu := \mathbb{E}_{y \sim P} \left[ \log \frac{P(y)}{Q(y)} \right] = \text{KL}(P \parallel Q) \quad \text{and} \quad \sigma = \sqrt{\text{Var}_{y \sim P} \left[ \log \frac{P(y)}{Q(y)} \right]}.$$



We first recall from our calculation in (119) that

$$\max \left\{ \frac{|Q(y) - P(y)|}{Q(y)}, \frac{|Q(y) - P(y)|}{P(y)} \right\} \leq \frac{\sqrt{2\mu}}{\min\{P(y), Q(y)\}} = O(\sqrt{\mu}),$$

where the last identity follows since  $P(y)$  and  $Q(y)$  are both bounded away from 0. Here and below, the notation  $f(\mu) = O(\sqrt{\mu})$  means  $|f(\mu)| \leq c_0\sqrt{\mu}$  for some *universal* constant  $c_0 > 0$ . This fact tells us that

$$\frac{P(y)}{Q(y)}, \frac{Q(y)}{P(y)} \in [1 \pm O(\sqrt{\mu})], \quad (123)$$

giving

$$\log \left( \frac{Q(y)}{P(y)} \right) = \log \left( 1 + \frac{Q(y) - P(y)}{P(y)} \right) = \frac{Q(y) - P(y)}{P(y)} + O \left( \frac{|Q(y) - P(y)|^2}{P^2(y)} \right) = (1 + O(\sqrt{\mu})) \frac{Q(y) - P(y)}{P(y)}.$$

All of this allows one to write

$$\begin{aligned} \sigma^2 &= \mathbb{E}_{y \sim P} \left[ \left( \log \frac{Q(y)}{P(y)} \right)^2 \right] - \mu^2 = (1 + O(\sqrt{\mu})) \mathbb{E}_{y \sim P} \left[ \left( \frac{Q(y) - P(y)}{P(y)} \right)^2 \right] - \mu^2 \\ &\stackrel{(a)}{=} (1 + O(\sqrt{\mu})) \mathbb{E}_{y \sim Q} \left[ \left( \frac{Q(y) - P(y)}{Q(y)} \right)^2 \right] - \mu^2 = (1 + O(\sqrt{\mu})) \chi^2(P \parallel Q) - \mu^2 \\ &= 2(1 + O(\sqrt{\mu}))\mu. \end{aligned}$$

Here, (a) arises due to (123), as the difference  $Q(y)/P(y)$  can be absorbed into the prefactor  $1 + O(\mu)$  by adjusting the constant in  $O(\mu)$  appropriately. The last line follows since, by [79, Proposition 2],

$$\frac{\mu}{\chi^2(P \parallel Q)} = \frac{\text{KL}(P \parallel Q)}{\chi^2(P \parallel Q)} = \left( 1 + O \left( \max_y \left\{ \max_y \frac{Q(y)}{P(y)}, \max_y \frac{P(y)}{Q(y)} \right\} \right) \right) \frac{1}{2} = (1 + O(\sqrt{\mu})) \frac{1}{2}.$$

Furthermore, it follows from [67, Fact 1] that

$$(4 - \log R) \text{H}^2(P, Q) \leq \text{KL}(P \parallel Q) \leq (4 + 2 \log R) \text{H}^2(P, Q),$$

where  $R := \max \left\{ \max_y \frac{P(y)}{Q(y)}, \max_y \frac{Q(y)}{P(y)} \right\}$ . In view of (123), it is seen that  $\log R = O(\sqrt{\mu})$ , thus establishing (51).

## E Proof of Lemma 4

It is tempting to use the matrix Bernstein inequality [80], but it loses a logarithmic factor in comparison to the desired bound given in Lemma 4. It turns out to be better to resort to Talagrand's inequality [81].

The starting point is to use the standard moment method and reduce to the case with independent entries (which has been studied in [82, 83]). Specifically, a standard symmetrization argument [84, Section 2.3] gives

$$\mathbb{E}[\|\mathbf{M}\|] \leq \sqrt{2\pi} \mathbb{E}[\|\mathbf{B}\|], \quad (124)$$

where  $\mathbf{B} = [\mathbf{B}_{i,j}]_{1 \leq i,j \leq n} := [g_{i,j} \mathbf{M}_{i,j}]_{1 \leq i,j \leq n}$  is obtained by inserting i.i.d. standard Gaussian variables  $\{g_{i,j} \mid i \geq j\}$  in front of  $\{\mathbf{M}_{i,j}\}$ . In order to upper bound  $\|\mathbf{B}\|$ , we further recognize that  $\|\mathbf{B}\|^{2p} \leq \text{Tr}(\mathbf{B}^{2p})$  for all  $p \in \mathbb{Z}$ . Expanding  $\mathbf{B}^{2p}$  as a sum over cycles of length  $2p$  and conditioning on  $\mathbf{M}$ , we have

$$\begin{aligned} \mathbb{E}[\text{Tr}(\mathbf{B}^{2p}) \mid \mathbf{M}] &= \sum_{1 \leq i_1, \dots, i_{2p} \leq n} \mathbb{E}[\text{Tr}(\mathbf{B}_{i_1, i_2} \mathbf{B}_{i_2, i_3} \cdots \mathbf{B}_{i_{2p-1}, i_{2p}} \mathbf{B}_{i_{2p}, i_1}) \mid \mathbf{M}] \\ &= \sum_{1 \leq i_1, \dots, i_{2p} \leq n} \text{Tr}(\mathbf{M}_{i_1, i_2} \cdots \mathbf{M}_{i_{2p}, i_1}) \mathbb{E} \left[ \prod_{j=1}^{2p} g_{i_j, i_{j+1}} \right] \end{aligned} \quad (125)$$

with the cyclic notation  $i_{2p+1} = i_1$ . The summands that are non-vanishing are those in which each distinct edge is visited an even number of times [84], and these summands obey  $\mathbb{E}[\prod_{j=1}^{2p} g_{i_j, i_{j+1}}] \geq 0$ . As a result,

$$(125) \leq \sum_{1 \leq i_1, \dots, i_{2p} \leq n} m \left( \prod_{j=1}^{2p} \|\mathbf{M}_{i_j, i_{j+1}}\| \right) \mathbb{E} \left[ \prod_{j=1}^{2p} g_{i_j, i_{j+1}} \right]. \quad (126)$$

We make the observation that the right-hand side of (126) is equal to  $m \cdot \mathbb{E}[\text{Tr}(\mathbf{Z}^{2p}) \mid \mathbf{M}]$ , where  $\mathbf{Z} := [g_{i,j} \|\mathbf{M}_{i,j}\|]_{1 \leq i,j \leq n}$ . Following the argument in [83, Section 2] and setting  $p = \log n$ , one derives

$$(\mathbb{E}[\text{Tr}(\mathbf{Z}^{2p}) \mid \mathbf{M}])^{\frac{1}{2p}} \lesssim \sigma + K\sqrt{2\log n},$$

where  $\sigma := \sqrt{\max_i \sum_j \|\mathbf{M}_{i,j}\|^2}$ . Putting all of this together, we obtain

$$\begin{aligned} \mathbb{E}[\|\mathbf{B}\| \mid \mathbf{M}] &\leq (\mathbb{E}[\|\mathbf{B}\|^{2p} \mid \mathbf{M}])^{\frac{1}{2p}} \leq (\mathbb{E}[\text{Tr}(\mathbf{B}^{2p}) \mid \mathbf{M}])^{\frac{1}{2p}} \\ &\leq m^{\frac{1}{2p}} \cdot (\mathbb{E}[\text{Tr}(\mathbf{Z}^{2p}) \mid \mathbf{M}])^{\frac{1}{2p}} \lesssim \sigma + K\sqrt{\log n}, \end{aligned} \quad (127)$$

where the last inequality follows since  $m^{\frac{1}{\log n}} \lesssim 1$  as long as  $m = n^{O(1)}$ . Combining (124) and (127) and undoing the conditional expectation yield

$$\mathbb{E}[\|\mathbf{M}\|] \leq \sqrt{2\pi} \mathbb{E}[\|\mathbf{B}\|] \lesssim \mathbb{E}[\sigma + K\sqrt{\log n}]. \quad (128)$$

Furthermore, Markov's inequality gives  $\mathbb{P}\{\|\mathbf{M}\| \geq 2\mathbb{E}[\|\mathbf{M}\|]\} \leq 1/2$  and hence

$$\text{Median}[\|\mathbf{M}\|] \leq 2\mathbb{E}[\|\mathbf{M}\|] \lesssim \mathbb{E}[\sigma + K\sqrt{\log n}]. \quad (129)$$

Now that we have controlled the expected spectral norm of  $\mathbf{M}$ , we can obtain concentration results by means of Talagrand's inequality [81]. See [84, Pages 73-75] for an introduction.

**Proposition 1 (Talagrand's inequality).** *Let  $\Omega = \Omega_1 \times \dots \times \Omega_N$  and  $\mathbb{P} = \mu_1 \times \dots \times \mu_N$  form a product probability measure. Each  $\Omega_l$  is equipped with a norm  $\|\cdot\|$ , and  $\sup_{x_l \in \Omega_l} \|x_l\| \leq K$  holds for all  $1 \leq l \leq N$ . Define  $d(x, y) := \sqrt{\sum_{l=1}^N \|x_l - y_l\|^2}$  for any  $x, y \in \Omega$ , and let  $f : \Omega \rightarrow \mathbb{R}$  be a 1-Lipschitz convex function with respect to  $d(\cdot, \cdot)$ . Then there exist some absolute constants  $C, c > 0$  such that*

$$\mathbb{P}\{|f(x) - \text{Median}[f(x)]| \geq \lambda K\} \leq C \exp(-c\lambda^2), \quad \forall \lambda. \quad (130)$$

Let  $\Omega_1, \dots, \Omega_N$  represent the sample spaces for  $\mathbf{M}_{1,1}, \dots, \mathbf{M}_{n,n}$ , respectively, and take  $\|\cdot\|$  to be the spectral norm. Clearly, for any  $\mathbf{M}, \tilde{\mathbf{M}} \in \mathbb{R}^{nm \times nm}$ , one has

$$\left| \|\mathbf{M}\| - \|\tilde{\mathbf{M}}\| \right| \leq \|\mathbf{M} - \tilde{\mathbf{M}}\| \leq \sqrt{\sum_{i \leq j} \|\mathbf{M}_{i,j} - \tilde{\mathbf{M}}_{i,j}\|^2} := d(\mathbf{M}, \tilde{\mathbf{M}}).$$

Consequently, Talagrand's inequality together with (129) implies that with probability  $1 - O(n^{-11})$ ,

$$\|\mathbf{M}\| \lesssim \text{Median}[\|\mathbf{M}\|] + K\sqrt{\log n} \lesssim \mathbb{E}[\sigma] + K\sqrt{\log n}. \quad (131)$$

Finally, if  $\mathbb{P}\{\mathbf{M}_{i,j} = \mathbf{0}\} = p_{\text{obs}}$  for some  $p_{\text{obs}} \gtrsim \log n/n$ , then the Chernoff bound when combined with the union bound indicates that

$$\sum_j \|\mathbf{M}_{i,j}\|^2 \leq K^2 \sum_j \mathbb{I}\{\mathbf{M}_{i,j} = \mathbf{0}\} \lesssim K^2 n p_{\text{obs}}, \quad 1 \leq i \leq n$$

with probability  $1 - O(n^{-10})$ , which in turn gives

$$\mathbb{E}[\sigma] \lesssim (1 - O(n^{-10})) \sqrt{K^2 n p_{\text{obs}}} + O(n^{-10}) K\sqrt{n} \lesssim K\sqrt{n p_{\text{obs}}}.$$

This together with (131) as well as the assumption  $p_{\text{obs}} \gtrsim \log n/n$  concludes the proof.

## F Proof of Lemma 5

The first step is to see that

$$\mathbf{L} - \mathbb{E}[\mathbf{L}] = \mathbf{L}^{\text{debias}} - \mathbb{E}[\mathbf{L}^{\text{debias}}],$$

where  $\mathbf{L}^{\text{debias}}$  is a debiased version of  $\mathbf{L}$  given in (12). This can be shown by recognizing that

$$\mathbf{1}^\top \mathbf{L}_{i,j} \mathbf{1} = \sum_{1 \leq \alpha, \beta \leq m} \ell(z_i = \alpha, z_j = \beta; y_{i,j}) = \sum_{1 \leq \alpha, \beta \leq m} \log \mathbb{P}(\eta_{i,j} = y_{i,j} - (\alpha - \beta)) = m \sum_{z=0}^{m-1} \log \mathbb{P}(\eta_{i,j} = z)$$

for any  $i \neq j$ , which is a fixed constant irrespective of  $y_{i,j}$ .

The main advantage to work with  $\mathbf{L}^{\text{debias}}$  is that each entry of  $\mathbf{L}^{\text{debias}}$  can be written as a linear combination of the log-likelihood ratios. Specifically, for any  $1 \leq \alpha, \beta \leq m$ ,

$$\begin{aligned} (\mathbf{L}_{i,j}^{\text{debias}})_{\alpha,\beta} &= \ell(z_i = \alpha, z_j = \beta; y_{i,j}) - \frac{1}{m} \sum_{l=0}^{m-1} \log P_l(y_{i,j}) = \log \mathbb{P}_{\alpha-\beta}(y_{i,j}) - \frac{1}{m} \sum_{l=0}^{m-1} \log P_{\alpha-\beta+l}(y_{i,j}) \\ &= \frac{1}{m} \sum_{l=1}^{m-1} \log \frac{P_{\alpha-\beta}(y_{i,j})}{P_{\alpha-\beta+l}(y_{i,j})} = \frac{1}{m} \sum_{l=1}^{m-1} \log \frac{P_0(y_{i,j} - \alpha + \beta)}{P_l(y_{i,j} - \alpha + \beta)}. \end{aligned} \quad (132)$$

Since  $\mathbf{L}_{i,j}^{\text{debias}}$  is circulant, its spectral norm is bounded by the  $\ell_1$  norm of any of its column:

$$\|\mathbf{L}_{i,j}^{\text{debias}}\| \leq \sum_{\alpha=1}^m |(\mathbf{L}_{i,j}^{\text{debias}})_{\alpha,1}| \leq \frac{1}{m} \sum_{l=1}^{m-1} \sum_{\alpha=1}^m \left| \log \frac{P_0(y_{i,j} - \alpha + 1)}{P_l(y_{i,j} - \alpha + 1)} \right| \quad (133)$$

$$= \frac{1}{m} \sum_{l=1}^{m-1} \sum_{y=1}^m \left| \log \frac{P_0(y)}{P_l(y)} \right| \leq \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1. \quad (134)$$

To finish up, apply Lemma 4 to arrive at

$$\|\mathbf{L} - \mathbb{E}[\mathbf{L}]\| = \|\mathbf{L}^{\text{debias}} - \mathbb{E}[\mathbf{L}^{\text{debias}}]\| \lesssim \left( \max_{i,j} \|\mathbf{L}_{i,j}^{\text{debias}}\| \right) \sqrt{np_{\text{obs}}} \lesssim \left( \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1 \right) \sqrt{np_{\text{obs}}}.$$

## G Proof of Lemma 6 and Lemma 8

(1) **Proof of Lemma 6.** In view of (65), for each  $1 \leq i \leq n$  one has

$$\|\mathbf{r}_i\|_\infty \leq \|\mathbb{E}[\mathbf{L}] \mathbf{h}\|_\infty + \|\mathbf{g}_i\|, \quad (135)$$

where  $\mathbf{g} := \tilde{\mathbf{L}} \mathbf{z}$ . In what follows, we will look at each term on the right-hand side of (135) separately.

- *The 1st term on the right-hand side of (135).* The feasibility constraint  $\mathbf{z}_i, \mathbf{x}_i \in \Delta$  implies  $\mathbf{1}^\top \mathbf{h}_j = 0$ , which enables us to express the  $i$ th block of  $\mathbf{f} := \frac{1}{p_{\text{obs}}} \mathbb{E}[\mathbf{L}] \mathbf{h}$  as

$$\mathbf{f}_i = \sum_{j:j \neq i} \mathbf{K} \mathbf{h}_j = \sum_{j:j \neq i} \mathbf{K}^0 \mathbf{h}_j, \quad ,$$

with  $\mathbf{K}^0$  defined in (35). By letting  $\mathbf{h}_{j,\setminus 1} := [h_{j,2}, \dots, h_{j,m}]^\top$  and denoting by  $\mathbf{K}_{:,l}$  (resp.  $\mathbf{K}_{l,:}$ ) the  $l$ th column (resp. row) of  $\mathbf{K}$ , we see that

$$\mathbf{f}_i = \mathbf{K}_{:,1}^0 \left( \sum_{j:j \neq i} h_{j,1} \right) + \mathbf{K}^0 \sum_{j:j \neq i} \begin{bmatrix} 0 \\ \mathbf{h}_{j,\setminus 1} \end{bmatrix}. \quad (136)$$

Recall from the feasibility constraint that  $h_{j,1} \leq 0$  and  $\mathbf{h}_{j,\setminus 1} \geq \mathbf{0}$ . Since  $\mathbf{K}^0$  is a non-positive matrix, one sees that the first term on the right-hand side of (136) is non-negative, whereas the second term is

non-positive. As a result, the  $l$ th entry of  $\mathbf{f}_i$ —denoted by  $f_{i,l}$ —is bounded in magnitude by

$$\begin{aligned} |f_{i,l}| &= \left| K_{l,1}^0 \left( \sum_{j:j \neq i} h_{j,1} \right) + \mathbf{K}_{l,:}^0 \sum_{j:j \neq i} \begin{bmatrix} 0 \\ \mathbf{h}_{j,\setminus 1} \end{bmatrix} \right| \\ &\leq \max \left\{ |K_{l,1}^0| \sum_{j:j \neq i} |h_{j,1}|, \left| \mathbf{K}_{l,:}^0 \cdot \sum_{j:j \neq i} \begin{bmatrix} 0 \\ \mathbf{h}_{j,\setminus 1} \end{bmatrix} \right| \right\} \\ &\leq \max \left\{ \text{KL}_{\max} \sum_{j:j \neq i} |h_{j,1}|, \text{KL}_{\max} \sum_{j:j \neq i} |\mathbf{1}^\top \mathbf{h}_{j,\setminus 1}| \right\} \\ &= \text{KL}_{\max} \sum_{j:j \neq i} |h_{j,1}|, \end{aligned}$$

where the last identity arises since  $\mathbf{1}^\top \mathbf{h}_{j,\setminus 1} = -h_{j,1}$ . Setting  $\bar{\mathbf{h}} = [\bar{h}_i]_{1 \leq i \leq m} := \frac{1}{n} \sum_{i=1}^n \mathbf{h}_i$  we obtain

$$\|\mathbf{f}_i\|_\infty \leq \text{KL}_{\max} \sum_{j=1}^n |h_{j,1}| = n \text{KL}_{\max} |\bar{h}_1| \leq n \text{KL}_{\max} \|\bar{\mathbf{h}}\|_\infty.$$

If we can further prove that

$$\|\bar{\mathbf{h}}\|_\infty \leq \min \left\{ \frac{k}{n}, \epsilon \right\}, \quad (137)$$

then we will arrive at the upper bound

$$\|\mathbb{E}[\mathbf{L}] \mathbf{h}\|_\infty = \|p_{\text{obs}} \mathbf{f}_i\|_\infty \leq n p_{\text{obs}} \text{KL}_{\max} \|\bar{\mathbf{h}}\|_\infty \leq n p_{\text{obs}} \text{KL}_{\max} \min \left\{ \frac{k}{n}, \epsilon \right\}. \quad (138)$$

To see why (137) holds, we observe that (i) the constraint  $\mathbf{z}_i \in \Delta$  implies  $|h_{i,l}| \leq 1$ , revealing that

$$\|\bar{\mathbf{h}}\|_\infty = \max_{1 \leq l \leq m} \frac{1}{n} \sum_{i=1}^n |h_{i,l}| \leq \frac{1}{n} \|\mathbf{h}\|_{*,0} = \frac{k}{n},$$

and (ii) by Cauchy-Schwarz,

$$\|\bar{\mathbf{h}}\|_\infty = \frac{1}{n} \left\| \begin{bmatrix} \mathbf{I}_m & \cdots & \mathbf{I}_m \end{bmatrix} \mathbf{h} \right\|_\infty \leq \frac{1}{n} (\sqrt{n} \cdot \|\mathbf{h}\|) = \epsilon.$$

- It remains to bound the 2nd term on the right-hand side of (135). Making use of Lemma 5 gives

$$\|\mathbf{g}\| \leq \|\tilde{\mathbf{L}}\| \|\mathbf{z}\| \lesssim \left\{ \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1 \right\} \sqrt{n p_{\text{obs}}} \sqrt{n} \quad (139)$$

with probability  $1 - O(n^{-10})$ . Let  $\|\mathbf{g}\|_{(1)} \geq \|\mathbf{g}\|_{(2)} \geq \cdots \geq \|\mathbf{g}\|_{(n)}$  denote the order statistics of  $\|\mathbf{g}_1\|, \dots, \|\mathbf{g}_n\|$ . Then, for any  $0 < \rho < 1$ ,

$$\|\mathbf{g}\|_{(\rho k^*)}^2 \leq \frac{1}{\rho k^*} \sum_{i=1}^{\rho k^*} \|\mathbf{g}\|_{(i)}^2 \leq \frac{1}{\rho k^*} \|\mathbf{g}\|^2 \lesssim \frac{n^2 p_{\text{obs}} \left\{ \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1 \right\}^2}{\rho k^*}, \quad (140)$$

where  $k^*$  is defined in (60). In addition, we have  $k^* \asymp n$  in the large-error regime (63). Substitution into (140) yields

$$\|\mathbf{g}\|_{(\rho k^*)} \lesssim \left\{ \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1 \right\} n \sqrt{\frac{p_{\text{obs}}}{\rho k^*}} \asymp \left\{ \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1 \right\} \sqrt{\frac{p_{\text{obs}} n}{\rho}}. \quad (141)$$

Consequently, if we denote by  $\mathcal{I}$  the index set of those blocks  $\mathbf{g}_i$  satisfying  $\|\mathbf{g}_i\| \lesssim \left\{ \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1 \right\} \sqrt{\frac{p_{\text{obs}} n}{\rho}}$ , then one has

$$|\mathcal{I}| \geq n - \rho k^*.$$

We are now ready to upper bound  $\|\mathbf{r}_i\|_\infty$ . For each  $i \in \mathcal{I}$ ,

$$\begin{aligned}
(135) &\leq np_{\text{obs}} \text{KL}_{\text{max}} \min \left\{ \frac{k}{n}, \epsilon \right\} + O \left( \left\{ \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1 \right\} \sqrt{\frac{np_{\text{obs}}}{\rho}} \right) \\
&\leq np_{\text{obs}} \text{KL}_{\text{max}} \min \left\{ \frac{k}{n}, \epsilon \right\} + \alpha np_{\text{obs}} \text{KL}_{\text{max}}
\end{aligned} \tag{142}$$

for some arbitrarily small constant  $\alpha > 0$ , provided that

$$np_{\text{obs}} \text{KL}_{\text{max}} \geq c_9 \left\{ \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1 \right\} \sqrt{\frac{p_{\text{obs}} n}{\rho}} \tag{143}$$

for some sufficiently large constant  $c_9 > 0$ . This can be satisfied if we pick  $\rho = c_5 \frac{\left\{ \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1 \right\}^2}{np_{\text{obs}} \text{KL}_{\text{max}}^2}$  for some sufficiently large constant  $c_5 > 0$ . Furthermore, in order to guarantee  $\rho < 1$ , one would need

$$\frac{\text{KL}_{\text{max}}^2}{\left\{ \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1 \right\}^2} \geq \frac{c_{10}}{p_{\text{obs}} n} \tag{144}$$

for some sufficiently large constant  $c_{10} > 0$ .

It is noteworthy that if  $m$  is fixed and if  $\min_{l,y} P_l(y)$  is bounded away from 0, then

$$\left\| \log \frac{P_0}{P_l} \right\|_1 \leq \sum_{y: P_0(y) \geq P_l(y)} \log \left( 1 + \frac{P_0(y) - P_l(y)}{P_l(y)} \right) + \sum_{y: P_0(y) < P_l(y)} \log \left( 1 + \frac{P_l(y) - P_0(y)}{P_0(y)} \right) \tag{145}$$

$$\leq \frac{1}{\min_{l,y} P_l(y)} \sum_y |P_0(y) - P_l(y)| \asymp \text{TV}(P_0, P_l) \tag{146}$$

$$\stackrel{(a)}{\lesssim} \sqrt{\text{KL}(P_0 \| P_l)} \leq \sqrt{\text{KL}_{\text{max}}}, \tag{147}$$

where (a) comes from Pinsker's inequality. Thus, in this case (144) would follow if  $\text{KL}_{\text{max}} \gg 1/(np_{\text{obs}})$ .

(2) **Proof of Lemma 8.** This part can be shown using similar argument as in the proof of Lemma 6. Specifically, from the definition (73) we have

$$\|\mathbf{q}_i\|_\infty \leq \|\mathbb{E}[\mathbf{L}] \mathbf{h}\|_\infty + \|\hat{\mathbf{g}}_i\| \leq np_{\text{obs}} \text{KL}_{\text{max}} \min \left\{ \frac{k}{n}, \epsilon \right\} + \|\hat{\mathbf{g}}_i\|, \tag{148}$$

where  $\hat{\mathbf{g}} := \tilde{\mathbf{L}} \mathbf{h}$ , and the last inequality is due to (138). Similar to (139) and (140), we get

$$\begin{aligned}
\|\hat{\mathbf{g}}\| &\leq \|\tilde{\mathbf{L}}\| \|\mathbf{h}\| \lesssim \left\{ \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1 \right\} \sqrt{np_{\text{obs}}} (\epsilon \sqrt{n}), \\
\|\hat{\mathbf{g}}\|_{(\rho k^*)} &\leq \frac{1}{\sqrt{\rho k^*}} \|\hat{\mathbf{g}}\| \asymp \frac{\left\{ \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1 \right\} (\epsilon \sqrt{n}) \sqrt{np_{\text{obs}}}}{\sqrt{\rho k^*}} \stackrel{(b)}{\lesssim} \left\{ \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1 \right\} \sqrt{\frac{np_{\text{obs}}}{\rho}},
\end{aligned}$$

where (b) arises since

$$\begin{aligned}
\frac{\epsilon \sqrt{n}}{\sqrt{k^*}} &= \frac{\|\mathbf{h}\|}{\sqrt{k}} \stackrel{(c)}{\leq} \frac{\sqrt{2\|\mathbf{h}\|_{*,0}}}{\sqrt{k}} \leq \sqrt{2}, \quad \text{if } k \leq \epsilon^2 n; \\
\frac{\epsilon \sqrt{n}}{\sqrt{k^*}} &= \frac{\epsilon \sqrt{n}}{\sqrt{\epsilon^2 n}} = 1, \quad \text{if } k > \epsilon^2 n.
\end{aligned}$$

Here, (c) results from Fact 1. One can thus find an index set  $\mathcal{I}$  with cardinality  $|\mathcal{I}| \geq n - \rho k^*$  such that

$$\|\hat{\mathbf{g}}_i\| \lesssim \left\{ \frac{1}{m} \sum_{l=1}^{m-1} \left\| \log \frac{P_0}{P_l} \right\|_1 \right\} \sqrt{\frac{np_{\text{obs}}}{\rho}}, \quad i \in \mathcal{I}, \tag{149}$$

where the right-hand side of (149) is identical to that of (141). Putting these bounds together and repeating the same argument as in (142)-(147) complete the proof.

## H Proof of Lemma 7

By definition, for each  $1 \leq i \leq n$  one has

$$s_{i,1} - s_{i,l} = \sum_{j: (i,j) \in \Omega} \log \frac{P_0(y_{i,j})}{P_{l-1}(y_{i,j})}, \quad 2 \leq l \leq m,$$

which is a sum of independent log-likelihood statistics. The main ingredient to control  $s_{i,1} - s_{i,l}$  is to establish the following lemma.

**Lemma 11.** *Consider two sequences of probability distributions  $\{P_i\}$  and  $\{Q_i\}$  on a finite set  $\mathcal{Y}$ . Generate  $n$  independent random variables  $y_i \sim P_i$ .*

(1) *For any  $\gamma \geq 0$ ,*

$$\mathbb{P} \left\{ \sum_{i=1}^n \log \frac{P_i(y_i)}{Q_i(y_i)} \leq \gamma \right\} \leq \exp \left\{ - \sum_{i=1}^n \mathbf{H}^2(P_i, Q_i) + \frac{1}{2} \gamma \right\}, \quad (150)$$

where  $\mathbf{H}^2(P_i, Q_i) := \frac{1}{2} \sum_y (\sqrt{P_i(y)} - \sqrt{Q_i(y)})^2$ .

(2) *Suppose that  $\left| \log \frac{Q_i(y_i)}{P_i(y_i)} \right| \leq K$ . If*

$$\min_{1 \leq i \leq n} \text{KL}(P_i \parallel Q_i) \geq c_2 \left\{ \frac{1}{n} \sqrt{\sum_{i=1}^n \text{Var}_{y_i \sim P_i} \left[ \log \frac{Q_i(y_i)}{P_i(y_i)} \right] \log(mn)} + \frac{K \log(mn)}{n} \right\}$$

for some sufficiently large constant  $c_2 > 0$ , then

$$\sum_{i=1}^n \log \frac{Q_i(y_i)}{P_i(y_i)} \leq -\frac{1}{2} n \min_{1 \leq i \leq n} \text{KL}(P_i \parallel Q_i) \quad (151)$$

with probability at least  $1 - O(m^{-11}n^{-11})$ .

We start from the case where  $m$  is fixed. When  $p_{\text{obs}} > c_0 \log n/n$  for some sufficiently large  $c_0 > 0$ , it follows from the Chernoff bound that

$$|\{j : (i, j) \in \Omega\}| \geq (1 - \zeta)np_{\text{obs}}, \quad 1 \leq i \leq n \quad (152)$$

with probability at least  $1 - O(n^{-10})$ , where  $\zeta > 0$  is some small constant. Taken together, Lemma 11(1) (with  $\gamma$  set to be  $2\zeta np_{\text{obs}} \mathbf{H}_{\min}^2$ ), (152), and the union bound give

$$s_{i,1} - s_{i,l} > 2\zeta np_{\text{obs}} \mathbf{H}_{\min}^2, \quad 1 \leq i \leq n, \quad 2 \leq l \leq m$$

or, equivalently,

$$\mathcal{S}(\mathbf{s}_i) > 2\zeta np_{\text{obs}} \mathbf{H}_{\min}^2, \quad 1 \leq i \leq n \quad (153)$$

with probability exceeding  $1 - mn \exp \{ - (1 - 2\zeta) np_{\text{obs}} \mathbf{H}_{\min}^2 \}$ . As a result, (153) would follow with probability at least  $1 - \exp \{ -\zeta \log(mn) \} - O(n^{-10})$ , as long as

$$\mathbf{H}_{\min}^2 \geq \frac{1 + \zeta}{1 - 2\zeta} \cdot \frac{\log n + \log m}{np_{\text{obs}}}. \quad (154)$$

It remains to translate these results into a version based on the KL divergence. Under Assumption 1, it comes from [67, Fact 1] that  $\mathbf{H}^2(P_0, P_l)$  and  $\text{KL}(P_0 \parallel P_l)$  are orderwise equivalent. This allows to rewrite (153) as

$$\mathcal{S}(\mathbf{s}_i) > c_5 \zeta np_{\text{obs}} \text{KL}_{\min}, \quad 1 \leq i \leq n \quad (155)$$

for some constant  $c_5 > 0$ . In addition, Lemma 3 and [67, Fact 1] reveal that

$$\mathbf{H}_{\min}^2 \geq \frac{1}{4} \left(1 - c_6 \sqrt{\mathbf{KL}_{\min}}\right) \mathbf{KL}_{\min} \quad \text{and} \quad \mathbf{H}_{\min}^2 \geq c_8 \mathbf{KL}_{\min}$$

for some constants  $c_6, c_8 > 0$ . As a result, (154) would hold if

$$\frac{(1 - c_6 \sqrt{\mathbf{KL}_{\min}}) \mathbf{KL}_{\min}}{4} \geq \frac{1 + \zeta}{1 - 2\zeta} \cdot \frac{\log n + \log m}{np_{\text{obs}}} \quad (156)$$

$$\text{or} \quad c_8 \mathbf{KL}_{\min} \geq \frac{1 + \zeta}{1 - 2\zeta} \cdot \frac{\log n + \log m}{np_{\text{obs}}} \quad (157)$$

When both  $\frac{\log n}{np_{\text{obs}}}$  and  $\zeta$  are sufficiently small, it is not hard to show that (156) is a consequence of  $\mathbf{KL}_{\min} \geq 4.01 \log n / (np_{\text{obs}})$ .

Finally, the second part (77) of Lemma 7 is straightforward by combining (152), Lemma 11(2), and the union bound.

**Proof of Lemma 11.** (1) For any  $\gamma \geq 0$ , taking the Chernoff bound we obtain

$$\mathbb{P} \left\{ \sum_{i=1}^n \log \frac{Q_i(y_i)}{P_i(y_i)} \geq -\gamma \right\} \leq \frac{\prod_{i=1}^n \mathbb{E}_{y \sim P_i} \left[ \exp \left( \frac{1}{2} \log \frac{Q_i(y)}{P_i(y)} \right) \right]}{\exp \left( -\frac{1}{2} \gamma \right)} = \exp \left( \frac{1}{2} \gamma \right) \prod_{i=1}^n (1 - \mathbf{H}^2(P_i, Q_i)),$$

where the last identity follows since

$$\begin{aligned} \mathbb{E}_{y \sim P} \left[ \exp \left( \frac{1}{2} \log \frac{Q(y)}{P(y)} \right) \right] &= \mathbb{E}_{y \sim P} \left[ \sqrt{\frac{Q(y)}{P(y)}} \right] = \sum_y \sqrt{P(y)Q(y)} \\ &= \sum_y \frac{1}{2} \left\{ P(y) + Q(y) - \left( \sqrt{P(y)} - \sqrt{Q(y)} \right)^2 \right\} = 1 - \mathbf{H}^2(P, Q). \end{aligned}$$

The claim (75) then follows by observing that  $1 - \mathbf{H}^2(P_i, Q_i) \leq \exp \{-\mathbf{H}^2(P_i, Q_i)\}$ .

(2) Taking expectation gives

$$\sum_{i=1}^n \mathbb{E}_{y_i \sim P_i} \left[ \log \frac{Q_i(y_i)}{P_i(y_i)} \right] = -\mathbf{KL}(P_i \parallel Q_i).$$

From our assumption  $\left| \log \frac{Q_i(y_i)}{P_i(y_i)} \right| \leq K$ , the Bernstein inequality ensures the existence of some constants  $c_0, c_1 > 0$  such that

$$\begin{aligned} &\sum_{i=1}^n \left( \log \frac{Q_i(y_i)}{P_i(y_i)} + \mathbf{KL}(P_i \parallel Q_i) \right) \\ &\leq c_0 \sqrt{\sum_{i=1}^n \text{Var}_{y_i \sim P_i} \left[ \log \frac{Q_i(y_i)}{P_i(y_i)} \right] \log(mn)} + c_1 \max_i \left| \log \frac{Q_i(y_i)}{P_i(y_i)} \right| \log(mn) \end{aligned} \quad (158)$$

with probability at least  $1 - O(m^{-11}n^{-11})$ . This taken collectively with the assumption

$$n \min_{1 \leq i \leq n} \mathbf{KL}(P_i \parallel Q_i) \gg \sqrt{\sum_{i=1}^n \text{Var}_{y_i \sim P_i} \left[ \log \frac{Q_i(y_i)}{P_i(y_i)} \right] \log(mn)} + K \log(mn)$$

establishes (151).  $\square$

## I Proof of Lemma 9

To begin with, (68) is an immediate consequence from (44) and Assumption 2. Next,

$$\frac{\text{KL}_{\max}}{\max_{l,y} \left| \log \frac{P_0(y)}{P_l(y)} \right|} \geq \frac{\text{KL}_{\max}}{\max_{1 \leq l < m} \left\| \log \frac{P_0}{P_l} \right\|_1} \stackrel{(a)}{\gg} \frac{1}{\sqrt{np_{\text{obs}}}} \stackrel{(b)}{\gtrsim} \frac{\log(mn)}{np_{\text{obs}}}.$$

where (a) arises from (44) together with Assumption 2, and (b) follows as soon as  $p_{\text{obs}} \gtrsim \frac{\log^2(mn)}{n}$ . This establishes the second property of (77).

Next, we turn to the first condition of (77). If  $\max_y P_0(y) \lesssim 1/\log(mn)$  holds, then we can derive

$$\begin{aligned} \text{Var}_{y \sim P_0} \left[ \log \frac{P_0(y)}{P_l(y)} \right] &\leq \sum_y P_0(y) \left( \log \frac{P_0(y)}{P_l(y)} \right)^2 \lesssim \frac{1}{\log(mn)} \sum_y \left( \log \frac{P_0(y)}{P_l(y)} \right)^2 \\ &\leq \frac{1}{\log(mn)} \left\| \log \frac{P_0}{P_l} \right\|_2^2, \end{aligned}$$

which together with (44) implies that

$$\frac{\text{KL}_{\min}^2}{\max_{0 \leq l < m} \text{Var}_{y \sim P_0} \left[ \log \frac{P_0(y)}{P_l(y)} \right]} \gtrsim \log(mn) \cdot \frac{\text{KL}_{\min}^2}{\max_{1 \leq l < m} \left\| \log \frac{P_0}{P_l} \right\|_1^2} \geq \frac{c_4 \log(mn)}{np_{\text{obs}}}. \quad (159)$$

Finally, consider the complement regime in which  $\max_y P_0(y) \gg \frac{1}{\log(mn)}$ , which obeys  $\max_y P_0(y) \gg 1/m$  as long as  $m \gtrsim \log(n)$ . Suppose  $P_0(y_0) = \max_y P_0(y)$  and  $\frac{P_0(y_0)}{P_l(y_0)} = \max_{j,y} \frac{P_0(y)}{P_j(y)}$  hold for some  $l$  and  $y_0$ , then it follows from the preceding inequality  $\max_y P_0(y) \gg 1/m$  that  $\frac{P_0(y_0)}{P_l(y_0)} \geq 2$ . Under Assumption 2, the KL divergence is lower bounded by

$$\text{KL}_{\min} \asymp \text{KL}_l \gtrsim H^2(P_0, P_l) \geq \frac{1}{2} \left( \sqrt{P_0(y_0)} - \sqrt{P_l(y_0)} \right)^2 \asymp P_0(y_0) \gg \frac{1}{\log(mn)}.$$

Moreover, the assumption  $m = n^{O(1)}$  taken collectively with Assumption 1 ensures

$$\max_{j,y} \left| \log \frac{P_0(y)}{P_j(y)} \right| \lesssim \log n,$$

allowing one to bound

$$\text{Var}_{y \sim P_0} \left[ \log \frac{P_0(y)}{P_l(y)} \right] \leq \max_{j,y} \left| \log \frac{P_0(y)}{P_j(y)} \right|^2 \lesssim \log^2 n.$$

Combining the above inequalities we obtain

$$\frac{\text{KL}_{\min}^2}{\max_{0 \leq l < m} \text{Var}_{y \sim P_0} \left[ \log \frac{P_0(y)}{P_l(y)} \right]} \gg \frac{\frac{1}{\log^2(mn)}}{\log^2 n} \gtrsim \frac{1}{\log^4(mn)} \gtrsim \frac{\log(mn)}{np_{\text{obs}}},$$

as long as  $p_{\text{obs}} \gtrsim \frac{\log^5(mn)}{n}$ , as claimed.

## J Proof of Lemma 10

For the sake of conciseness, we will only prove (100), as (101) can be shown using the same argument. We recognize that (100) can be established by demonstrating the moderate deviation principle with respect to

$$S_n = \sum_{i=1}^n (\zeta_{i,n} - \mathbb{E}[\zeta_{i,n}]) \quad \text{where } \zeta_{i,n} := \frac{\log \frac{Q_n(y_{i,n})}{P_n(y_{i,n})}}{\sigma_n}.$$



To be precise, the main step is to invoke [85, Theorem 6] to deduce that

$$a_n \log \mathbb{P} \left\{ \sqrt{a_n} \frac{S_n}{\sqrt{n}} > \tau \right\} = -(1 + o_n(1)) \frac{\tau^2}{2} \quad (160)$$

for any constant  $\tau > 0$ , where

$$a_n := \sigma_n^2 / (n\mu_n^2). \quad (161)$$

In fact, one can connect the event  $\left\{ \sqrt{a_n} \frac{S_n}{\sqrt{n}} > \tau \right\}$  with the likelihood ratio test because

$$\mathbb{P} \left\{ \sqrt{a_n} \frac{S_n}{\sqrt{n}} > \tau \right\} = \mathbb{P} \left\{ \sum_{i=1}^n \frac{\log \frac{Q_n(y_{i,n})}{P_n(y_{i,n})} + \mu_n}{\sigma_n} > n \frac{\mu_n}{\sigma_n} \tau \right\} = \mathbb{P} \left\{ \sum_{i=1}^n \log \frac{Q_n(y_{i,n})}{P_n(y_{i,n})} > (\tau - 1) n \mu_n \right\}.$$

This reveals that (160) is equivalent to

$$\mathbb{P} \left\{ \sum_{i=1}^n \log \frac{Q_n(y_{i,n})}{P_n(y_{i,n})} > (\tau - 1) n \mu_n \right\} = \exp \left( - (1 + o_n(1)) n \frac{\tau^2 \mu_n^2}{2 \sigma_n^2} \right) \quad (162)$$

as claimed in the first identity of (100). Moreover, by Lemma 3, it is seen that  $2\mu_n = (1 + O(\sqrt{\mu_n})) \sigma_n^2$  in the regime considered herein, leading to the second identity of (100). Hence, it suffices to prove (160).

In order to apply [85, Theorem 6], we need to check that the double indexed sequence  $\{\zeta_{i,n} - \mathbb{E}[\zeta_{i,n}] : i \leq n\}_{n \geq 1}$  satisfies the conditions required therein. First of all, the independence assumption gives

$$v^2 := \sup_{i,n} \left\{ \text{Var}[\zeta_{i,n}] + 2 \sum_{j>i} |\text{Cov}(\zeta_{i,n}, \zeta_{j,n})| \right\} = \sup_{i,n} \{\text{Var}[\zeta_{i,n}]\} = 1 < \infty.$$

Second,

$$\frac{\text{Var}[S_n]}{n} = \frac{\sum_{i=1}^n \text{Var}[\zeta_{i,n}]}{n} = 1 > 0.$$

Third, it follows from Lemma 2 that

$$M_n := \sup_i |\zeta_{i,n}| = \frac{\sup_i \left| \log \frac{Q_n(y_{i,n})}{P_n(y_{i,n})} \right|}{\sigma_n} \leq \frac{1}{\min\{P_n(y_{i,n}), Q_n(y_{i,n})\}} \frac{\sqrt{2\mu_n}}{\sigma_n} \asymp \frac{\sqrt{\mu_n}}{\sigma_n}. \quad (163)$$

The assumption  $\frac{\mu_n^2}{\sigma_n^2} \asymp \frac{\log n}{n}$  further gives

$$a_n = \frac{\sigma_n^2}{n\mu_n^2} \asymp \frac{1}{\log n} = o_n(1).$$

Moreover, making use of the bound (163) as well as the assumptions  $\frac{\mu_n^2}{\sigma_n^2} \asymp \frac{\log n}{n}$  and  $\mu_n \gtrsim \frac{\log n}{n}$ , we derive

$$\frac{na_n}{M_n^2 \log^4 n} \gtrsim \frac{n \frac{\sigma_n^2}{n\mu_n^2}}{\frac{\mu_n}{\sigma_n^2} \cdot \log^4 n} = \frac{\sigma_n^4}{\mu_n^3 \log^4 n} = \frac{\sigma_n^4}{\mu_n^4} \cdot \frac{\mu_n}{\log^4 n} \gtrsim \frac{n}{\log^5 n} \rightarrow \infty.$$

With these conditions in place, we can invoke [85, Theorem 6] to establish (160).

## References

- [1] R. Keshavan, A. Montanari, and S. Oh, “Matrix completion from a few entries,” *IEEE Transactions on Information Theory*, vol. 56, no. 6, pp. 2980–2998, 2010.
- [2] —, “Matrix completion from noisy entries,” *Journal of Machine Learning Research*, vol. 99, pp. 2057–2078, 2010.

- [3] P. Jain, P. Netrapalli, and S. Sanghavi, “Low-rank matrix completion using alternating minimization,” in *ACM symposium on Theory of computing*. ACM, 2013, pp. 665–674.
- [4] R. Sun and Z.-Q. Luo, “Guaranteed matrix completion via nonconvex factorization,” in *Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2015, pp. 270–289.
- [5] Y. Chen and M. J. Wainwright, “Fast low-rank estimation by projected gradient descent: General statistical and algorithmic guarantees,” *arXiv preprint arXiv:1509.03025*, 2015.
- [6] R. Ge, J. D. Lee, and T. Ma, “Matrix completion has no spurious local minimum,” *arXiv preprint arXiv:1605.07272*, 2016.
- [7] D. Park, A. Kyrillidis, S. Bhojanapalli, C. Caramanis, and S. Sanghavi, “Provable non-convex projected gradient descent for a class of constrained matrix optimization problems,” *arXiv preprint arXiv:1606.01316*, 2016.
- [8] Q. Zheng and J. Lafferty, “A convergent gradient descent algorithm for rank minimization and semidefinite programming from random linear measurements,” in *NIPS*, 2015, pp. 109–117.
- [9] S. Tu, R. Boczar, M. Soltanolkotabi, and B. Recht, “Low-rank solutions of linear matrix equations via procrustes flow,” *arXiv preprint arXiv:1507.03566*, 2015.
- [10] S. Bhojanapalli, B. Neyshabur, and N. Srebro, “Global optimality of local search for low rank matrix recovery,” *arXiv preprint arXiv:1605.07221*, 2016.
- [11] P. Netrapalli, P. Jain, and S. Sanghavi, “Phase retrieval using alternating minimization,” in *Advances in Neural Information Processing Systems*, 2013, pp. 2796–2804.
- [12] E. Candes, X. Li, and M. Soltanolkotabi, “Phase retrieval via wirtinger flow: Theory and algorithms,” *Information Theory, IEEE Transactions on*, vol. 61, no. 4, pp. 1985–2007, 2015.
- [13] Y. Shechtman, A. Beck, and Y. C. Eldar, “GESPAR: Efficient phase retrieval of sparse signals,” *IEEE transactions on signal processing*, vol. 62, no. 4, pp. 928–938, 2014.
- [14] Y. Chen and E. Candes, “Solving random quadratic systems of equations is nearly as easy as solving linear systems,” *accepted to Communications on Pure and Applied Mathematics*, 2015.
- [15] T. Cai, X. Li, and Z. Ma, “Optimal rates of convergence for noisy sparse phase retrieval via thresholded Wirtinger flow,” *arXiv preprint arXiv:1506.03382*, 2015.
- [16] J. Sun, Q. Qu, and J. Wright, “A geometric analysis of phase retrieval,” *arXiv:1602.06664*, 2016.
- [17] Y. Chi and Y. M. Lu, “Kaczmarz method for solving quadratic equations,” *IEEE Signal Processing Letters*, vol. 23, no. 9, pp. 1183–1187, 2016.
- [18] H. Zhang, Y. Chi, and Y. Liang, “Provable non-convex phase retrieval with outliers: Median truncated Wirtinger flow,” *International conference on machine learning*, June 2016.
- [19] J. Sun, Q. Qu, and J. Wright, “Complete dictionary recovery using nonconvex optimization,” in *Proceedings of the 32nd International Conference on Machine Learning*, 2015, pp. 2351–2360.
- [20] —, “When are nonconvex problems not scary?” *arXiv preprint arXiv:1510.06096*, 2015.
- [21] X. Li, S. Ling, T. Strohmer, and K. Wei, “Rapid, robust, and reliable blind deconvolution via nonconvex optimization,” *arXiv preprint arXiv:1606.04933*, 2016.
- [22] S. Mei, Y. Bai, and A. Montanari, “The landscape of empirical risk for non-convex losses,” *arXiv preprint arXiv:1607.06534*, 2016.
- [23] J. D. Lee, M. Simchowitz, M. I. Jordan, and B. Recht, “Gradient descent only converges to minimizers,” in *Conference on Learning Theory*, 2016, pp. 1246–1257.

- [24] B. Karrer and M. E. Newman, “Stochastic blockmodels and community structure in networks,” *Physical Review E*, vol. 83, no. 1, p. 016107, 2011.
- [25] A. Coja-Oghlan, “Graph partitioning via adaptive spectral techniques,” *Combinatorics, Probability and Computing*, vol. 19, no. 02, pp. 227–284, 2010.
- [26] E. Abbe, A. Bandeira, and G. Hall, “Exact recovery in the stochastic block model,” *IEEE Transactions on Information Theory*, vol. 62, no. 1, pp. 471–487, 2016.
- [27] K. Chaudhuri, F. C. Graham, and A. Tsiatas, “Spectral clustering of graphs with general degrees in the extended planted partition model,” *Journal of Machine Learning Research*, vol. 23, pp. 35–1, 2012.
- [28] E. Mossel, J. Neeman, and A. Sly, “Consistency thresholds for binary symmetric block models,” *Arxiv preprint*, 2014.
- [29] —, “Belief propagation, robust reconstruction and optimal recovery of block models,” in *COLT*, vol. 35, 2014, pp. 356–370.
- [30] A. Jalali, Y. Chen, S. Sanghavi, and H. Xu, “Clustering partially observed graphs via convex optimization,” *International Conf. on Machine Learning (ICML)*, 2011.
- [31] S. Oymak and B. Hassibi, “Finding dense clusters via "low rank+ sparse" decomposition,” *arXiv preprint arXiv:1104.5186*, 2011.
- [32] Y. Chen, S. Sanghavi, and H. Xu, “Clustering sparse graphs,” *NIPS*, December 2012.
- [33] B. Hajek, Y. Wu, and J. Xu, “Achieving exact cluster recovery threshold via semidefinite programming,” *IEEE Transactions on Information Theory*, vol. 62, no. 5, pp. 2788–2797, 2016.
- [34] —, “Achieving exact cluster recovery threshold via semidefinite programming: Extensions,” *arXiv preprint arXiv:1502.07738*, 2015.
- [35] E. Abbe, A. S. Bandeira, A. Bracher, and A. Singer, “Decoding binary node labels from censored edge measurements: Phase transition and efficient recovery,” *IEEE Transactions on Network Science and Engineering*, vol. 1, no. 1, pp. 10–22, 2014.
- [36] Y. Chen, G. Kamath, C. Suh, and D. Tse, “Community recovery in graphs with locality,” *International Conference on Machine Learning*, June 2016.
- [37] A. Javanmard, A. Montanari, and F. Ricci-Tersenghi, “Phase transitions in semidefinite relaxations,” *Proceedings of the National Academy of Sciences*, vol. 113, no. 16, pp. E2218–E2223, 2016.
- [38] P. Chin, A. Rao, and V. Vu, “Stochastic block model and community detection in the sparse graphs: A spectral algorithm with optimal rate of recovery,” *arXiv preprint arXiv:1501.05021*, vol. 2, no. 4, 2015.
- [39] A. Globerson, T. Roughgarden, D. Sontag, and C. Yildirim, “How hard is inference for structured prediction?” in *International Conference on Machine Learning*, 2015, pp. 2181–2190.
- [40] D. Hernando, P. Kellman, J. Haldar, and Z.-P. Liang, “Robust water/fat separation in the presence of large field inhomogeneities using a graph cut algorithm,” *Magnetic resonance in medicine*, vol. 63, no. 1, pp. 79–90, 2010.
- [41] T. Zhang, Y. Chen, S. Bao, M. Alley, J. Pauly, B. Hargreaves, and S. Vasanawala, “Resolving phase ambiguity in dual-echo dixon imaging using a projected power method,” *to appear in Magnetic Resonance in Medicine*, 2016.
- [42] J. Berglund, H. Ahlström, L. Johansson, and J. Kullberg, “Two-point dixon method with flexible echo times,” *Magnetic resonance in medicine*, vol. 65, no. 4, pp. 994–1004, 2011.
- [43] Q.-X. Huang, H. Su, and L. Guibas, “Fine-grained semi-supervised labeling of large shape collections,” *ACM Transactions on Graphics*, vol. 32, no. 6, p. 190, 2013.

- [44] L. Wang and A. Singer, “Exact and stable recovery of rotations for robust synchronization,” *Information and Inference*, 2013.
- [45] A. S. Bandeira, M. Charikar, A. Singer, and A. Zhu, “Multireference alignment using semidefinite programming,” in *Conference on Innovations in Theoretical Computer Science*, 2014, pp. 459–470.
- [46] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra, “Efficient projections onto the  $\ell_1$ -ball for learning in high dimensions,” in *International conference on Machine learning*, 2008, pp. 272–279.
- [47] G. H. Golub and C. F. Van Loan, *Matrix computations*. JHU Press, 2012, vol. 3.
- [48] M. Journée, Y. Nesterov, P. Richtárik, and R. Sepulchre, “Generalized power method for sparse principal component analysis,” *Journal of Machine Learning Research*, vol. 11, no. Feb, pp. 517–553, 2010.
- [49] M. Hein and T. Bühler, “An inverse power method for nonlinear eigenproblems with applications in 1-spectral clustering and sparse pca,” in *Neural Information Processing Systems*, 2010, pp. 847–855.
- [50] X.-T. Yuan and T. Zhang, “Truncated power method for sparse eigenvalue problems,” *Journal of Machine Learning Research*, vol. 14, no. Apr, pp. 899–925, 2013.
- [51] Y. Deshpande and A. Montanari, “Finding hidden cliques of size  $\sqrt{N/e}$  in nearly linear time,” *Foundations of Computational Mathematics*, vol. 15, no. 4, pp. 1069–1128, 2015.
- [52] N. Boumal, “Nonconvex phase synchronization,” *arXiv preprint arXiv:1601.06114*, 2016.
- [53] L. Massoulié, “Community detection thresholds and the weak Ramanujan property,” in *Symposium on Theory of Computing*. ACM, 2014, pp. 694–703.
- [54] E. Abbe and C. Sandon, “Community detection in general stochastic block models: fundamental limits and efficient recovery algorithms,” *arXiv preprint arXiv:1503.00609*, 2015.
- [55] G. Brito, I. Dumitriu, S. Ganguly, C. Hoffman, and L. V. Tran, “Recovery and rigidity in a regular stochastic block model,” in *Symposium on Discrete Algorithms*, 2016, pp. 1589–1601.
- [56] O. Guédon and R. Vershynin, “Community detection in sparse networks via Grothendieck’s inequality,” *Probability Theory and Related Fields*, pp. 1–25, 2015.
- [57] A. Singer, “Angular synchronization by eigenvectors and semidefinite programming,” *Applied and computational harmonic analysis*, vol. 30, no. 1, pp. 20–36, 2011.
- [58] A. Bandeira, N. Boumal, and A. Singer, “Tightness of the maximum likelihood semidefinite relaxation for angular synchronization,” *arXiv preprint arXiv:1411.3272*, 2014.
- [59] H. Liu, M.-C. Yue, and A. M.-C. So, “On the statistical performance of the generalized power method for angular synchronization,” *arXiv preprint arXiv:1603.00211*, 2016.
- [60] A. Bandeira, N. Boumal, and V. Voroninski, “On the low-rank approach for semidefinite programs arising in synchronization and community detection,” *arXiv preprint arXiv:1602.04426*, 2016.
- [61] E. Candès, X. Li, Y. Ma, and J. Wright, “Robust principal component analysis?” *Journal of ACM*, vol. 58, no. 3, pp. 11:1–11:37, Jun 2011.
- [62] V. Chandrasekaran, S. Sanghavi, P. Parrilo, and A. Willsky, “Rank-sparsity incoherence for matrix decomposition,” *SIAM Journal on Optimization*, vol. 21, no. 2, 2011.
- [63] A. Ganesh, J. Wright, X. Li, E. Candes, and Y. Ma, “Dense error correction for low-rank matrices via principal component pursuit,” in *International Symposium on Information Theory*, 2010, pp. 1513–1517.
- [64] Y. Chen, A. Jalali, S. Sanghavi, and C. Caramanis, “Low-rank matrix recovery from errors and erasures,” *IEEE Transactions on Information Theory*, vol. 59, no. 7, pp. 4324–4337, 2013.

- [65] Q. Huang and L. Guibas, “Consistent shape maps via semidefinite programming,” *Computer Graphics Forum*, vol. 32, no. 5, pp. 177–186, 2013.
- [66] Y. Chen, L. J. Guibas, and Q.-X. Huang, “Near-optimal joint object matching via convex relaxation,” *International Conference on Machine Learning (ICML)*, pp. 100 – 108, 2014.
- [67] Y. Chen, C. Suh, and A. J. Goldsmith, “Information recovery from pairwise measurements,” *accepted to IEEE Transactions on Information Theory*, *arXiv:1504.01369*, 2015.
- [68] A. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu, “ShapeNet: An information-rich 3D model repository,” *arXiv preprint arXiv:1512.03012*, 2015.
- [69] V. G. Kim, W. Li, N. J. Mitra, S. DiVerdi, and T. Funkhouser, “Exploring collections of 3d models using fuzzy correspondences,” *ACM Transactions on Graphics (TOG)*, vol. 31, no. 4, p. 54, 2012.
- [70] D. Pachauri, R. Kondor, and V. Singh, “Solving the multi-way matching problem by permutation synchronization.” in *Advanced in Neural Information Processing Systems (NIPS)*, 2013.
- [71] Y. Shen, Q. Huang, N. Srebro, and S. Sanghavi, “Normalized spectral map synchronization,” in *Neural Information Processing Systems (NIPS)*, 2016.
- [72] R. Jonker and A. Volgenant, “A shortest augmenting path algorithm for dense and sparse linear assignment problems,” *Computing*, vol. 38, no. 4, pp. 325–340, 1987.
- [73] D. Giorgi, S. Biasotti, and L. Paraboschi, “Shape retrieval contest 2007: Watertight models track,” *SHREC competition*, vol. 8, 2007.
- [74] A. Blake, P. Kohli, and C. Rother, *Markov random fields for vision and image processing*. Mit Press, 2011.
- [75] P. Ravikumar and J. Lafferty, “Quadratic programming relaxations for metric labeling and Markov random field MAP estimation,” *International Conference on Machine Learning*, pp. 737–744, 2006.
- [76] Q. Huang, Y. Chen, and L. Guibas, “Scalable semidefinite relaxation for maximum a posterior estimation,” *International Conference on Machine Learning*, 2014.
- [77] A. B. Tsybakov, *Introduction to nonparametric estimation*. Springer Science & Business Media, 2008.
- [78] I. Sason and S. Verdú, “f-divergence inequalities,” *arxiv preprint arXiv:1508.00335*, 2015.
- [79] S. Dragomir, “Upper and lower bounds for Csiszár’s f-divergence in terms of the Kullback-Leibler distance and applications,” *Inequalities for the Csiszár’s f-divergence in Information Theory*, 2000.
- [80] J. A. Tropp, “An introduction to matrix concentration inequalities,” *to appear in Foundations and Trends in Machine Learning*, 2015.
- [81] M. Talagrand, “Concentration of measure and isoperimetric inequalities in product spaces,” *Publications Mathématiques de l’Institut des Hautes Etudes Scientifiques*, vol. 81, no. 1, pp. 73–205, 1995.
- [82] Y. Seginer, “The expected norm of random matrices,” *Combinatorics, Probability and Computing*, vol. 9, no. 02, pp. 149–166, 2000.
- [83] A. Bandeira and R. van Handel, “Sharp nonasymptotic bounds on the norm of random matrices with independent entries,” *arXiv preprint arXiv:1408.6185*, 2014.
- [84] T. Tao, *Topics in random matrix theory*. AMS Bookstore, 2012, vol. 132.
- [85] F. Merlevède, M. Peligrad, E. Rio *et al.*, “Bernstein inequality and moderate deviations under strong mixing conditions,” in *High dimensional probability V: the Luminy volume*. Institute of Mathematical Statistics, 2009, pp. 273–292.